arXiv:1409.5414v4 [cs.DS] 19 Apr 2015

Differentially Private Linear Algebra in the Streaming Model

Jalaj Upadhyay
Center for Applied Cryptographic Research
David R. Cheriton School of Computer Science
University of Waterloo.
jalaj.upadhyay@uwaterloo.ca

Abstract

Numerical linear algebra plays an important role in computer science. In this paper, we initiate the study of performing linear algebraic tasks while preserving privacy when the data is streamed online. Our main focus is the space requirement of the privacy-preserving data-structures. We give the first *sketch-based* algorithm for differential privacy. We give optimal, up to logarithmic factor, space data-structures that can compute low rank approximation, linear regression, and matrix multiplication, while preserving differential privacy with better additive error bounds compared to the known results. Notably, we match the best known space bound in the non-private setting by Kane and Nelson (J. ACM, 61(1):4).

Our mechanism for differentially private low-rank approximation *reuses* the random Gaussian matrix in a specific way to provide a single-pass mechanism. We prove that the resulting distribution also preserve differential privacy. This can be of independent interest. We do not make any assumptions, like singular value separation or normalized row assumption, as made in the earlier works. The mechanisms for matrix multiplication and linear regression can be seen as the private analogues of the known non-private algorithms. All our mechanisms, in the form presented, can also be computed in the distributed setting.

Keywords. Differential Privacy, Linear Algebra, Random Projection.

1 Introduction

An $n \times d$ real-valued matrix is a natural structure for storing data about n entities described by d features. They arise in many contexts like information retrieval, data analysis, finance, scientific computation, genetics, and machine learning. In most of these applications, one is often required to do various linear algebraic tasks, like low-rank matrix approximations (LRA), linear regression (LIN-REG), and matrix multiplication (MAT-MULT). For example, LRA is used in latent semantic indexing to speed-up the computation by computing a new representation for each document in the original collection, LIN-REG is used in finance to analyze and quantify the systematic risk of an investment, and many matrix operations in scientific computation such as Gaussian elimination, LU decomposition, and the determinant or the inverse of a matrix can be reduced to MAT-MULT.

Let us consider some applications of these linear algebraic tasks in learning theory. For example, LRA is an effective tool in reducing the memory burdens of kernel methods. In kernel methods, we map the data into a kernel-induced Hilbert space. Given n samples, this requires the calculation of an $n \times n$ symmetric, positive semi-definite kernel matrix, which requires quadratic space and mostly cubic time. This can be very demanding for large problems, impeding its practical deployment. This is where LRA is helpful. Williams and Seeger [42] showed that the spectra of kernel matrices decay rapidly. This led to the following approach: compute the LRA of the kernel matrices and perform computation on those matrices. This simultaneously brought improvement in many other models, like support vector machine, kernel Fisher discriminant analysis and kernel principal component analysis. Similarly, LIN-REG, which is used to predict the output for a new data based on the previous data, has played pivotal role in machine learning. Examples include, *capital asset pricing model* which is used to predict demands, supplies, and investment, *Reinforcement Learning* to approximate value functions, and learning parameters of a model with linear dynamics (see, for eg., Walsh *et al.* [41]). Likewise, MAT-MULT has found many roles in machine learning other than modelling multivariate regression as a transposed MAT-MULT problem. For example, Mossel *et al.* [31] showed how to improve run time of learning juntas using matrix multiplication.

Motivation of this work. The data in the examples mentioned above and on which we perform the learning contain sensitive information and takes a lot of space. As a result, privacy and memory constraints are as important as correct computation. The privacy issue has been recently exemplified by the deanonymization of Netflix datasets, which was cited as one of the reasons to study differentially private low-rank approximation by [22, 23]. Similarly, data used in genetics engineering and finance have natural privacy concerns. This raises the question of whether one can perform all these tasks while giving a robust guarantee of privacy, like *differential privacy*. When privacy is not a concern, there are many algorithms which use almost optimal space and one-pass over the input matrix (see, for example, Kane and Nelson [27]). On the other hand, all the known private algorithms that achieve small additive error use multiple passes over the matrix. For example, the traditional *Krylov subspace iteration* method (on which some of the recent works like [21, 23, 28] are based) requires a lot of space and multiple passes over the input matrix. *The main focus of this work is to perform linear algebraic computation while preserving privacy of the data with the additional challenge that the data is received online and we can use sub-linear space.*

First sketch based differentially private mechanisms. A natural approach to perform linear algebraic tasks on online data while preserving privacy is to generate and store a sub-linear space data structure that can be used to perform these tasks without leaking privacy. In the non-private setting, the standard techniques either perform *random sampling* or generate a *random sketch* of the input data. Dwork *et al.* [14] showed that known (non-private) random sampling based algorithms for certain statistical queries can be made differentially private. They also gave an impossibility result for private analogues of *sketch based approaches* for specific "statistical queries." This raises doubts over the applicability of sketch based approach in privacy. In this paper, *we show the first set of positive results for sketch based approach on streamed*

data. We give almost optimal space data-structures that can compute LRA*,* LIN-REG*, and* MAT-MULT *while preserving differential privacy.* Notably, we match the best known space bounds for all these tasks in the non-private setting. We remark that our results does not contradict the impossibility result of Dwork *et al.* [14], which holds for cleverly constructed statistical queries.

PRIVACY MODEL USED IN THIS PAPER. There are two commonly used notions of differential privacy: *event level privacy*, where guarantees are at the granularity of individual records in the datasets, and *user level privacy*, where guarantees are at the granularity of each user whose data is present in the dataset. Dwork *et al.* [13] showed that it is impossible to obtain any non-trivial result with respect to the user level privacy on a streamed data. Therefore, in this paper, we use event level privacy. More specifically, we say two matrices $\mathbf{A}_1$ and $\mathbf{A}_2$ with same rows and columns are *neighbouring* if the matrix $\mathbf{A}_1 - \mathbf{A}_2$ is a rank-1 matrix and $\|\mathbf{A}_1 - \mathbf{A}_2\| \leq 1$. This notion was also used recently by [22, 23, 21] for differentially private low-rank approximation and by Blocki *et al.* [4, 5] and Upadhyay [39] for certain statistical queries. We also restrict how the mechanism receives the data in the sense that the data matrix is streamed online and arrives either row-wise or column-wise. We call two data streams neighbouring if they are formed by streaming entries of neighbouring matrices. With this notion of neighbouring datasets, we use the following definition for privacy.

Definition 1.1. A randomized algorithm $\mathfrak{M}$ gives (ε, δ)-*differential privacy*, if for all neighbouring data-streams $\mathbf{A}_1, \mathbf{A}_2$ and all $S \subset Range(\mathfrak{M})$, $\mathsf{Pr}[\mathfrak{M}(\mathbf{A}_1) \in S] \leq \exp(\varepsilon)\mathsf{Pr}[\mathfrak{M}(\mathbf{A}_2) \in S] + \delta$, where the probability is over the coin tosses of $\mathfrak{M}$.

The main focus of this study is the space requirement of the privacy-preserving data-structures.

1.1 PROBLEM STATEMENTS AND OUR RESULTS

In this section, we give the formal description of the problems we investigate in this paper. The performance of a streaming algorithm is measured by three basic factors: the number of passes over the data stream, the space used by the data-structure, and the time taken to update the data-structure. All our private mechanisms for performing linear algebraic tasks are single-pass and achieve almost optimal space bound for one-pass algorithms. Our private sketch are linear; therefore, *our mechanisms extends naturally to turnstile updates*. We reserve the letter n for the number of rows and d for the number of columns of a private matrix. We assume $d < n$. For *bit complexity*, we use the convention of Clarkson and Woodruff [10], i.e., the entries of a matrix can be represented by $\kappa = \log(nd)$ bits.

LOW-RANK APPROXIMATION. We start with the problem statement of LRA.

Problem 1. *((α, β, τ)-*LRA*).* Given an $n \times d$ matrix $\mathbf{A}$ and a target rank k, construct a matrix $\mathbf{\Psi}$ with k orthonormal columns such that $\|\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^{\mathsf{T}}\mathbf{A}\|_N \leq (1+\alpha)\min_{\mathsf{rank}(\mathbf{A}_k)\leq k} \|\mathbf{A} - \mathbf{A}_k\|_N + \tau$ with probability at least $1 - \beta$ where N represents either Frobenius or spectral norm.

The best known space lower bound in the non-private setting for $(\alpha, 5/6, 0)$-LRA is $\mathbf{\Omega}(nk\alpha^{-1}\kappa)$ due to Clarkson and Woodruff [10]. Kane and Nelson [27, Th 6.2] improved the earlier analysis of Clarkson and Woodruff [10] to show a one-pass non-private streaming algorithm for approximate LRA with respect to the Frobenius norm with row/column-wise updates where one maintains a data-structure in the form of a sketch of size $O(k\alpha^{-1}(n+d)\kappa\log(1/\beta))$.

The first approach to prove differential privacy using the sketch generation algorithm of Clarkson and Woodruff [10] is to use additive noise mechanisms. There are two main drawback of this approach. A straightforward analysis along the line of [22] with no coherence assumption leads to an additive error, τ, of order $k^{3/2}$ for the Frobenius norm approximation. The second drawback is that it only leads to approximation when the approximation norm is Frobenius. In this paper, we use random Gaussian matrices to generate the sketch and prove the following in Section 4 (see Theorem 4.1 for the precise statement).

Method	Additive noise (τ)	Privacy Notion	# Passes
Chaudhary *et al.* [9]	$O(nk/\varepsilon)$	Spectral norm	k
Hardt and Roth [22]	$\frac{\sqrt{kn}\log(k/\delta)}{\varepsilon} + \sqrt{\frac{\mu\|A\|_F \log(k/\delta)}{\varepsilon}}$	Event level	2
Hardt and Roth [23]	$O(\frac{k^2}{\varepsilon}\sqrt{(\mathsf{rk}(A)\mu + k\log n)\log\left(\frac{1}{\delta}\right)\log n})$	Event level	$k\sqrt{\log\sigma}$
Kapralov and Talwar [28]	$O(dk^3/(\varepsilon\gamma^2))$	Spectral norm	k
Hardt and Price [21]	$\frac{\sigma_1\sqrt{kn\mu\log(1/\delta)\log(n/\gamma)\log\log(n/\gamma)}}{\varepsilon\gamma^{1.5}\sigma_k}$	Event level	$k\sqrt{\log\sigma}$
Dwork *et al.* [17]	$O((k\sqrt{n}\ln(1/\delta))/\varepsilon) + \tilde{O}(\sqrt{k^3n^{3/2}}/\varepsilon^2)$	User level	1
This paper (Theorem 4.1(i))	$O(\sqrt{nk\ln(k/\delta)}/\varepsilon)$	Event level	1
This paper (Theorem 4.1(ii))	$O(\sqrt{nk\ln(k/\delta)}/\varepsilon)$	Event level	1

Table 1: Comparison Between our Mechanism and Previous Mechanisms for Differentially Private k-Rank Approximation of an $n \times d$ rank-$\mathsf{rk}(A)$ matrix A. μ denotes the coherence of $\mathbf{A}$, $\gamma = (\sigma_k/\sigma_{k+1}) - 1$. [22] and Theorem 4.1(i) bounds the Frobenius norm while the other entries bounds the Spectral norm.

Theorem (Informal). There is an $O(k\varepsilon^{-1}(n+d)\kappa)$ bits data structure that could be used to publish (ε,δ)-differentially private k-rank approximation of an $n \times d$ private matrix in a single pass with $\tau \leq O(\sqrt{nk\ln(2/\delta)}/\varepsilon)$ for the Frobenius and $\tau \leq O(\sqrt{nk\ln(2/\delta)}/\varepsilon)$ for the spectral norm.

Note that our data-structure provides LRA in both the spectral and Frobenius norm with the same space bound as in the non-private setting. The reason behind this is simple. The bound on the single pass algorithm of [10, 27] and the two-pass algorithm of Sarlos [36] use their bound for MAT-MULT (see Problem 3), and, therefore, only achieves a bound on the Frobenius norm. On the other hand, we use perturbation theory which gives us a unified bound for the both the norms. Notably, our proof does not rely on our bound in MAT-MULT (see our result for Problem 3).

COMPARISON AND TIGHTNESS. Let $\sigma_1, \cdots, \sigma_{\mathsf{rank}(\mathbf{A})}$ be the singular values of a matrix $\mathbf{A}$. We compare our result for LRA with the previous works in Table 1. Due to lack of space, we defer the detail comparison to Appendix 4.0.1. Here we just state few main points. Works like [9, 17, 28] consider different notion of neighbouring datasets, so our results are incomparable to theirs. Apart from the result in this paper, only Dwork *et al.* [17] give LRA on streamed data though under extra assumptions like *normalized row* assumption and a lower bound on the optimal value. Also, we do not make any *low-coherence* assumption. Therefore, if we set $\mu = n$ (for worst-case) to compare our results, then Table 1 shows that we achieve best bound on τ with the same privacy model. In fact, we achieve almost optimal τ for spectral norm (see, [23, Thm. 6.1]). Our result for Frobenius norm is also tight due to the result of Blum *et al.* [6].

LINEAR REGRESSION. We start by giving the problem statement.

Problem 2. *((α,β,τ)-LIN-REG).* Given an $n \times d$ matrix $\mathbf{A}$ and a m set of $n \times 1$ column vectors $\mathbf{B} = \{\mathbf{b}_1, \cdots, \mathbf{b}_m\}$, output a set of vectors $\mathbf{X} = \{\mathbf{x}_1, \cdots \mathbf{x}_m\}$ all in $\mathbb{R}^d$ so that for all $i \in [m]$, $\|\mathbf{A}\mathbf{x}_i - \mathbf{b}_i\|_F \leq (1+\alpha)\min_{\mathbf{y}\in\mathbb{R}^{d\times 1}} \|\mathbf{A}\mathbf{y} - \mathbf{b}_i\|_F + \tau$ with probability at least $1 - \beta/m$.

The best known space lower bound for $(\alpha, 5/6, 0)$-LIN-REG in the non-private setting is $\Omega(d^2\alpha^{-1}\kappa)$ due to [10]. Kane and Nelson [27, Th 6.2] improved the analysis of Clarkson and Woodruff [10] to gave a $(d^2\alpha^{-1}\kappa\log(1/\beta))$ bits data-structure that can answer $(\alpha,\beta,0)$-LIN-REG in the non-private setting. We show that their space bound can be achieved in the private setting at the expense of a small additive error. Informally, we show the following in Section 5.2 (see Theorem 5.3 for the precise statement).

Theorem (Informal). There is a data-structure that uses $O\left(d^2\alpha^{-1}\kappa\log\left(1/\beta\right)\right)$ bits of space to compute $(\alpha,\beta,O(s^2\alpha\sqrt{n}))$-LIN-REG while providing (ε,δ)-differential privacy, where $s =$

Problem	Lower Bound	Non-private Algorithm	This paper
(α, β, τ)-MAT-MULT	$\Omega\left(d\alpha^{-2}\kappa\right)$	$\tilde{O}(d\alpha^{-2}\kappa\log(1/\beta))$	$\tilde{O}(d\alpha^{-2}\kappa\log(1/\beta))$
(α, β, τ)-LIN-REG	$\Omega\left(d^2\alpha^{-1}\kappa\right)$	$\tilde{O}\left(d^2\alpha^{-1}\kappa\log(1/\beta)\right)$	$\tilde{O}\left(d^2\alpha^{-1}\kappa\log(1/\beta)\right)$
Frobenius norm (α, β, τ)-LRA	$\Omega(nk\kappa/\alpha)$	$\tilde{O}(k\alpha^{-1}(n+d)\kappa)$	$O(k\varepsilon^{-1}(n+d)\kappa)$
Spectral norm (α, β, τ)-LRA	–	–	$O(k\varepsilon^{-1}(n+d)\kappa)$

Table 2: Our Results with Respect to the Best Known Space Bounds in Non-private Setting.

$\sqrt{\frac{16d\log(1/\beta)\ln(2/\delta)}{\alpha\varepsilon^2}}\ln\left(\frac{16d\log(1/\beta)}{\alpha\delta}\right)$.

LIN-REG has been also studied in the *local privacy model* by Duchi *et al.* [11] and in the *online private learning model* by Jain *et al.* [25] and Thakurata and Smith [38]. These models are different from ours; therefore, our results are incomparable to theirs.

MATRIX MULTIPLICATION. We start by giving the problem statement.

Problem 3. *((α, β, τ)-MAT-MULT).* An $n \times d_1$ matrix $\mathbf{A}$ and $n \times d_2$ matrix $\mathbf{B}$ are given. Output a matrix $\mathbf{C}$ so that $\|\mathbf{A}^\mathsf{T}\mathbf{B} - \mathbf{C}\|_F \leq \alpha\|\mathbf{A}\|_F \cdot \|\mathbf{B}\|_F + \tau$ with probability at least $1 - \beta$.

Let $d = \max\{d_1, d_2\}$. The best known space lower bound for $(\alpha, 5/6, 0)$-MAT-MULT in the non-private setting is $\Omega(d\alpha^{-2}\kappa)$ due to Clarkson and Woodruff [10]. Kane and Nelson [27, Th 6.2] improved the analysis of Clarkson and Woodruff [10] to give a $O\left(d\alpha^{-2}\kappa\log\left(1/\beta\right)\right)$-bits data-structure that can solve $(\alpha, \beta, 0)$-MAT-MULT in the non-private setting. In Section 5.1, we show that this bound can be achieved in the private setting with a small additive error (see Theorem 5.1 for the precise statement).

> **Theorem (Informal).** Given conforming matrices $\mathbf{A}$ and $\mathbf{B}$. Let $r = O(\log(1/\beta)/\alpha^2)$ and $s = \sqrt{16r\ln(2/\delta)}\varepsilon^{-1}\ln(16r/\delta)$. For large enough $d = \max\{d_1, d_2\}$, there is data structure which maintains a sketch of size $O\left(d\alpha^{-2}\kappa\log\left(1/\beta\right)\right)$, computes $(\alpha, \beta, O(s^2\alpha\sqrt{n}))$-MAT-MULT and provides (ε, δ)-differential privacy.

SPACE OPTIMALITY OF THE DATA STRUCTURES. In Table 2, we give the best known lower bound for the space required for all the problems we study in this paper, that by the best known data-structures in the non-private setting, and by our private data-structures. The parameters used in Table 2 are the same as in Table 1. As one can see, we achieve almost optimal space data-structures for all the three problems. Note that, we explicitly require $\mathbf{\Omega}$ to be stored for the mechanism of LRA; in the other two cases, $\mathbf{\Omega}$ can be picked from the distribution on the fly. We note that stronger lower bounds on space are achievable for each of the problems considered because of the non-zero additive error; however, the bounds stated in the table are presently the best known lower bounds that we can use to make any sort of optimality comparison.

One can also implement our mechanisms *as distributed algorithms*, a desirable feature as argued by [3]. This is because all the operations used in our mechanisms have efficient distributed algorithms (Jacobi method for singular value decomposition, Cannon [8]'s algorithm for multiplication, and GMRES [35]'s residual method).

1.2 OUR TECHNIQUES

We start by giving an overview of the generic private sketch algorithm and then discuss the ideas used in this work which are different from other related works.

PRIVATE-SKETCH GENERATION. In order to get a better bound on the additive noise, we devise a private-sketch generation (PSG) mechanism to generate a private sketch of the input stream. We reduce the privacy

guarantee to maintaining a spectral property of the streamed matrix. At a high level, all our mechanisms use this basic mechanism while maintaining the spectral properties of the private matrix. We note that the affine transformation of Blocki *et al.* [4, 5] and Upadhyay [39] could be also used to maintain the required spectral properties, but this leads to a large additive error. Moreover, as noted by Blocki *et al.* [4], it does not guarantee LRA. For these reasons, we use a different method.

MECHANISM AND ANALYSIS FOR LRA. We follow two-steps to compute a LRA: (i) compute a projection matrix (*range finding* step), (ii) compute a LRA by operating the projection matrix on the input matrix (*projection* step). In the most naive form, both the steps require the input matrix, and hence two-passes are required. The first observation is that the mechanism for PSG already gives considerable improvement in the range finding step. This by itself does not give LRA as the singular values of the projection matrix and private matrix are not comparable at this stage. We need to use the projection step. The second observation is that, by a clever use of linear algebra, information gathered in the range-finding step using a random Gaussian matrix can be used to emulate the projection step without using the input matrix explicitly. However, we need to reuse the random Gaussian matrix. This complicates the privacy proof (discussed below in more detail). We also use the idea of Hardt and Price [21] to use an oversampling parameter p to get a sharper bounds for both the Frobenius and spectral norm.

DIFFERENCES IN TECHNIQUES TO COMPUTE LRA. Our mechanism for LRA is markedly different from the recent private mechanisms [9, 17, 21, 22, 23, 28]. Among all these mechanisms, only Dwork *et al.* [17] computes LRA privately in an online manner by using binary tree technique of Dwork *et al.* [13]. We differ at a basic level from these mechanisms: all these works perturb the output by adding noise to it, while *we perturb the input matrix and then multiply noise matrix*. Our mechanism uses random Gaussian matrices (privacy proof holds only for Gaussian matrix); therefore, we use different tricks, like reusing Gaussian matrix, as mentioned in the last paragraph.

DIFFERENCES FROM PREVIOUS ANALYSES OF LRA. Our analysis has some interesting features. As mentioned above, we reuse the Gaussian matrix to get a mechanism that uses only one-pass over the input matrix. In general, reusing randomness can result in privacy breach. Fortunately, the Gaussian matrix is reused in a specific manner for which we prove that privacy holds under certain spectral property of the input matrix. We believe this can be of independent interest. Our mechanism works by maintaining the required singular value for the input matrix. This causes the additive error. The analysis used in this paper to bound the error differs a lot from the analyses of Clarkson and Woodruff [10] and Sarlos [36]. They use the trick that a good bound on MAT-MULT allows a bound on LRA. This limits its applicability to the case when the approximation metric is the Frobenius norm. Moreover, Sarlos [36] use two-passes over the input matrix. We use perturbation theory, which allows us to give bounds for both the norms in an unified manner. We believe this can be of independent interest.

OTHER RESULTS: MECHANISM FOR MAT-MULT AND LIN-REG. At a high level, our mechanisms for (α, β, τ)-MAT-MULT and (α, β, τ)-LIN-REG are private analogues of Clarkson and Woodruff [10]. Their algorithm use *tug-of-war matrices* with Rademacher entries (this helps in improving the update time). Adding Gaussian noise to ensure differential privacy amounts to a large additive error. We use PSG mechanism. For the privacy proof to go through, we need to lift the singular values of the input matrix. Our choice to lift the singular values is constrained by keeping a check on τ as well as to keep the mechanism one-pass. In order to use the proofs of Kane and Nelson [27] to give an approximation bound, we use an analogous variance bound on random Gaussian matrix (since privacy does not necessarily hold for random matrices used by Kane and Nelson [27]).

RELATED WORKS. The first formal definition and mechanism for differential privacy was given by Dwork *et al.* [12]. Since then, many mechanisms for preserving differential privacy have been proposed

in the literature (see, Dwork and Roth [15]). The literature on non-private streaming algorithms is so extensive that we cannot hope to cover it in any detail here. In the private setting, Dwork *et al.* [14] studied *pan-privacy*, where the internal state is known to the adversary. Subsequently, there have been some works on online private learning, like Dwork *et al.* [17], Jain *et al.* [25], and Thakurata and Smith [38], for various tasks. There are some recent works on differentially private LRA as well. Blum *et al.* [6] first studied this problem and gave a simple "input perturbation" algorithm that adds noise to the covariance matrix. This was improved by Hardt and Roth [22] who studied LRA for the Frobenius norm under the low coherence assumption. Kapralov and Talwar [28] and Chaudhary *et al.* [9] studied the spectral LRA of a matrix by giving a matching upper and lower bounds for privately computing the top k eigenvectors of a matrix. Hardt and Roth [23] and Hardt and Price [21] improved their noise bound by proposing robust private subspace iteration mechanism. Recently, Dwork *et al.* [17] revisited randomized mechanism to give a tighter bound and used it to give an private online learning algorithm under a normalized row assumption.

2 Notations and Basic Preliminaries

We reserve the letters $\mathbf{A}$ and $\mathbf{B}$ for private input matrices, and $\mathbf{\Omega}$ for a random Gaussian matrix. For an $n \times d$ matrix $\mathbf{A}$, we let $\mathbf{A}_{i:}$ denote the i-th row of $\mathbf{A}$, $\mathbf{A}_{:j}$ denote the j-th column of $\mathbf{A}$, and $\mathbf{A}'$ denote the symmetric matrix $\begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^\mathsf{T} & 0 \end{pmatrix}$ corresponding to $\mathbf{A}$. We let $\mathbf{A}_t$ denote the matrix received after t time epochs. The singular value decomposition (SVD) of $\mathbf{A}$ is $\mathbf{A} = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^\mathsf{T}$, where $\mathbf{U}$ and $\mathbf{V}$ are left and right eigenvectors of $\mathbf{A}$, and $\mathbf{\Sigma}$ is a diagonal matrix. The entries of Λ are called the *singular values* of $\mathbf{A}$. Since $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, one can write $\mathbf{A}^i = \mathbf{V}\mathbf{\Sigma}^i\mathbf{U}^\mathsf{T}$ for any real value i. We let $\mathsf{rank}(\mathbf{A})$ denotes the rank of the matrix $\mathbf{A}$ and $\sigma_i(\mathbf{A})$ its singular values. Where it is clear from context, we write σ_i for the singular values.

We use various matrix norm. We use the notation $\|\cdot\|_F$ for Frobenius norm. A Frobenius norm for a matrix $\mathbf{A} = (a_{ij})_{i\in[n],j\in[d]}$ is defined as following $\|\mathbf{A}\|_F = \sum_{ij} |a_{ij}|^2$. For a matrix $\mathbf{A}$, we let $\|\mathbf{A}\|_2$ denote the 2-norm, i.e., $\max_{\mathbf{x}\in\mathbb{R}^d} \|\mathbf{A}\mathbf{x}\|_2/\|\mathbf{x}\|_2$. s When we wish to refer to both the Frobenius and the spectral norm, we overload the symbol $\|\cdot\|$ and drop the subscript. We let $\mathbf{e}_1, \cdots, \mathbf{e}_d$ denote the standard basis vectors in $\mathbb{R}^d$. We use bold face symbols to denote vectors and $\mathbf{0}^n$ to denote an n-dimensional 0-vector. For a matrix $\mathbf{M}$, we write $\mathbf{M} \succeq 0$ if all its eigenvalues are non-zero. In the course of this paper, we use many standard results from linear algebra and random matrices.

LINEAR ALGEBRA. Our analysis make extensive use of linear algebra and statistical properties of Gaussian distribution. We give an exposition to the level required to understand this paper.

Lemma 2.1. If matrix $\mathbf{A}$ and $\mathbf{B}$ are conforming, then $\|\mathbf{A}\mathbf{B}\|_F \leq \|\mathbf{A}\|_2\|\mathbf{B}\|_F$.

Lemma 2.2. Let $\mathbf{A}$ and $\mathbf{B}$ be Hermittian matrices. If $\mathbf{A}$ and $\mathbf{B}$ differ in at most one row of Euclidean norm 1, then $\mathrm{Tr}(\mathbf{A}^\mathsf{T}\mathbf{A}) - \mathrm{Tr}(\mathbf{B}^\mathsf{T}\mathbf{B}) \leq 2$.

UNIVARIATE AND MULTIVARIATE GAUSSIAN DISTRIBUTION. A random variable, X, distributed according to a Gaussian distribution has the probability density function, $\mathsf{PDF}_X(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$. We denote it by $X \sim \mathcal{N}(\mu, \sigma^2)$. The Gaussian distribution is invariant under affine transformation, i.e., if $X \sim \mathcal{N}(\mu_x, \sigma_x)$ and $Y \sim \mathcal{N}(\mu_y, \sigma_y)$, then $Z = aX + bY$ has the distribution $Z \sim \mathcal{N}(a\mu_x + b\mu_y, a\sigma_x^2 + b\sigma_y^2)$.

The multivariate Gaussian distribution is a generalization of univariate Gaussian distribution. Given a m dimensional multivariate random variable, $X \sim \mathcal{N}(\mu, \mathbf{\Sigma})$, the PDF of a multivariate Gaussian is given by $\mathsf{PDF}_\mathbf{X}(\mathbf{x}) := \frac{1}{\sqrt{(2\pi)^{\mathsf{rank}(\mathbf{\Sigma})}\mathsf{Det}(\mathbf{\Sigma})}} \exp\left(-\frac{1}{2}\mathbf{x}^\mathsf{T}\mathbf{\Sigma}^\dagger\mathbf{x}\right)$ with mean $\mu \in \mathbb{R}^m$ and covariance matrix $\mathbf{\Sigma} = \mathbb{E}[(X-\mu)(X-\mu)^\mathsf{T}]$. If $\mathbf{\Sigma}$ has a non-trivial kernel space, then the PDF is undefined. However, in this paper, we only need to compare the probability distribution of two random variables which are defined over the

On input a streamed column $\mathbf{v} \in \mathbb{R}^{2\tilde{n}}$ of the private matrix, parameters $r, \tilde{n}$, the mechanism samples an $r \times 2\tilde{n}$ random matrix $\mathbf{\Omega}$, and

PSG_1**:** Compute $\mathbf{Y_v} = \mathbf{\Omega v}$, and return $\mathbf{Y_v}$.

PSG_2**:** Compute $\mathbf{Y_v} = \mathbf{\Omega}^\mathsf{T}\mathbf{\Omega v}$, and return $\mathbf{Y_v}$.

Figure 1: Private Sketch Generation (PSG) Algorithm

same subspace. Therefore, wherever required, we restrict our attention to the (sub)space orthogonal to the kernel space of $\mathbf{\Sigma}$. Multivariate Gaussian distribution maintains many key properties of univariate Gaussian distribution. For example, any (non-empty) subset of multivariate Gaussians is a multivariate Gaussian and linear functions of multivariate Gaussian random variables are multivariate Gaussian random variables, i.e., if $\mathbf{y} = \mathbf{Ax} + \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{n\times n}$ is a non-singular matrix and $\mathbf{b} \in \mathbb{R}^n$, then $\mathbf{y} \sim \mathcal{N}(\mathbf{A}\mu + \mathbf{b}, \mathbf{A\Sigma A}^\mathsf{T})$.

We use the following properties of Gaussian matrices, a proof of the first lemma could be found in [37], while for the rest two, one can refer to [32].

Lemma 2.3. Let $\mathbf{\Omega}$ be a Gaussian matrix. Then $\mathbb{E}[\|\mathbf{A\Omega B}\|_2] \leq \|\mathbf{A}\|_2\|\mathbf{B}\|_F + \|\mathbf{A}\|_F\|\mathbf{B}\|_2$ for fixed matrices $\mathbf{A}$ and $\mathbf{B}$.

Lemma 2.4. Let $\mathbf{\Omega}$ be a $n \times (k+p)$ Gaussian matrix. Then

$$\mathbb{E}[\|\mathbf{\Omega}^{-1}\|_F^2] = \sqrt{\frac{k}{p-1}} \qquad \text{and} \qquad \mathbb{E}[\|\mathbf{\Omega}^{-1}\|_2] \leq \frac{e\sqrt{k+p}}{p}.$$

Lemma 2.5. Let $\mathbf{\Omega}$ be a random $n \times (k+p)$ Gaussian matrix whose entries are picked from the distribution $\mathcal{N}(0,1)$. Then $\mathbb{E}[\mathrm{Tr}((\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega})^{-1})] = k/(p-1)$.

Theorem 2.6. *(Johnson-Lindenstrauss lemma)* Fix any $\eta < 1/2$ and let m be a positive integer. Let $\mathbf{\Omega}$ be a $k \times n$ matrix with entries picked from a Gaussian distribution $\mathcal{N}(0,1)$, where $k \geq 4(\alpha^2/2 - \alpha^3/3)^{-1} \ln m$. Then for any m unit vector set S in $\mathbb{R}^n$

$$\forall \mathbf{x} \in S, \mathsf{Pr}_\mathbf{M}\left[\|\mathbf{\Omega x}\|_2 \in (1 \pm \alpha)\|\mathbf{x}\|_2\right] \geq 1 - 2\exp(-\alpha^2 k/8).$$

DIFFERENTIAL PRIVACY. We use the following in our analysis explicitly or implicitly.

Theorem 2.7. *([16]).* Let $\varepsilon, \delta \in (0,1)$, and $\delta' > 0$. If $\mathcal{K}_1, \cdots, \mathcal{K}_\ell$ are each (ε, δ)-differential private mechanism, then the mechanism $\mathcal{K}(D) := (\mathcal{K}_1(D), \cdots, \mathcal{K}_\ell(D))$ releasing the concatenation of each algorithm is $(\varepsilon', \ell\delta + \delta')$-differentially private for $\varepsilon' < \sqrt{2\ell \ln(1/\delta')}\varepsilon + 2\ell\varepsilon^2$.

Lemma 2.8. Let $M(D)$ be a (ε, δ)-differential private mechanism for a database D , and let h be any function, then any mechanism $M' := h(M(D))$ is also (ε, δ)-differentially private for the same set of queries.

3 Differentially Private Sketch Generation

We study differential privacy in the well known streaming model of computation [1]. We present it at the level required to understand this paper (a more formal definition can be found in [1]). This model has three entities: a stream generator $\mathcal{S}$, a (database) curator $\mathcal{K}$, and a query maker $\mathcal{Q}$. $\mathcal{S}$ starts the process at time $t = 0$ and the curator initializes its data structure to $\mathcal{D}_0$. Thereafter, the curator is allowed only one-pass over the input matrix, i.e., it can access any entry of the data-base during exactly one time epoch, and update its data structure to $\mathcal{D}_t$ using $\mathcal{D}_{t-1}$ and the newly accessed data-points of the matrix. At any time, t, the query maker $\mathcal{Q}$ makes a query q. The curator responds with $q(\mathcal{D}_t)$. The streaming model has a resource bound on

the curator. A curator is only allowed to use total time *polynomial in the size of the data-base* to construct the data structure, where the size of data-structure should be *sub-linear in the size of the data-base*. For *differential privacy, we further require that the response of $\mathcal{K}$ to the query of $\mathcal{Q}$ should satisfy Definition 1.1 with respect to two neighbouring streams.* Differential privacy on streaming data has been also studied as private online learning by [25] and [38], where the emphasis is on learning about the streamed input – one bears a regret on a hypothesis evaluated against a data point which is not yet streamed. Due to lack of space, we do not delve into the detail comparison of the two models. Our main privacy result is as follows.

Theorem 3.1. (i) If the singular values of the matrix whose columns are streamed to PSG_1 are at least $\sigma_{\min} := \left(4\sqrt{r\log(2/\delta)}\log(r/\delta)\right)/\varepsilon$, then PSG_1 preserves (ε,δ)-differential privacy. (ii) If the singular values of the matrix whose columns are streamed to PSG_2 algorithm are at least $\sigma_{\min} := (4r\log(r/\delta))/\varepsilon$, then PSG_2 preserves (ε,δ)-differential privacy.

The proof of part (i) needs some care due to subtleties mentioned later in this section, but the overall outline follows the idea of Blocki *et al.* [4]. However, the proof of part (ii) is very different and more involved. We show that for a streamed matrix $\mathbf{A}$, the PDF of the published matrix using PSG_2 is

$$\frac{\exp(-\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{\Phi})/2)\Delta(\mathbf{\Phi})^{(n-r-1)/2}}{2^{rn/2}\pi^{r(r-1)/4}\Delta(\mathbf{A})^{r/2}\prod_{i=1}^{r}\Gamma((n-i+1)/2)}, \tag{1}$$

where $\Delta(\cdot)$ is the determinant and $\mathbf{\Phi} = \sum_{i=1}^{r}\mathbf{a}_i\mathbf{a}_i^{\mathsf{T}}$, where $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}^n, \mathbb{I}_{n\times n})$. To achieve equation (2), we use the chain rule to break the computation to smaller part. We then compute the probability distribution function of every term in the chain rule. Note that equation (2) is not a Wishart distribution as $\mathbf{A}^{-1}\Phi$ need not be symmetric. This is the most technical part of the proof and we perform this computation from the basics; rest of the proof requires clever manipulation of the PDF for neighbouring matrices.

PSG_2 *is interesting in its own regard. It says that two successive applications of a Gaussian matrix in a defined form preserves differential privacy*. In general, reuse of randomness does not preserve privacy, but what we show here is that if the randomness is reused cleverly, then it is possible to achieve privacy.

Proof. We start with the proof of part (i). We give the proof [4] for the sake of completion. We denote by $\widetilde{\mathbf{A}}$ the matrix that differs from $\mathbf{A}$ by at most one entry. In other word, if $\mathbf{A}$ and $\widetilde{\mathbf{A}}$ differs in row i, then there exists a unit vector v such that $\mathbf{A} - \widetilde{\mathbf{A}} = \mathbf{E} = \mathbf{v}\mathbf{e}_i^{\mathsf{T}}$. Let $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathsf{T}}$ ($\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{V}}^{\mathsf{T}}$, respectively) be the SVD of $\mathbf{A}$ ($\widetilde{\mathbf{A}}$, respectively).

The PDF for the two distributions, corresponding to $\mathbf{A}$ and $\widetilde{\mathbf{A}}$, is just a linear transformation of $\mathcal{N}(0,\mathbb{I}_{n\times n})$. Therefore,

$$\mathsf{PDF}_{\mathbf{A}^{\mathsf{T}}Y}(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^d\Delta(\mathbf{A}^{\mathsf{T}}\mathbf{A})}}\exp(-\frac{1}{2}\mathbf{x}(\mathbf{A}^{\mathsf{T}}\mathbf{A})^{-1}\mathbf{x}^{\mathsf{T}})$$
$$\mathsf{PDF}_{\widetilde{\mathbf{A}}^{\mathsf{T}}Y}(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^d\Delta(\widetilde{\mathbf{A}}^{\mathsf{T}}\widetilde{\mathbf{A}})}}\exp(-\frac{1}{2}\mathbf{x}(\widetilde{\mathbf{A}}^{\mathsf{T}}\widetilde{\mathbf{A}})^{-1}\mathbf{x}^{\mathsf{T}})$$

We prove the result for a row of the published matrix; the theorem follows from Theorem 2.7. It is straightforward to see that combination of the following proves differential privacy for a row of published matrix:

$$\sqrt{\frac{\Delta(\mathbf{A}^{\mathsf{T}}\mathbf{A})}{\Delta(\widetilde{\mathbf{A}}^{\mathsf{T}}\widetilde{\mathbf{A}})}} \in \exp(\pm\varepsilon_0) \quad \text{and} \quad \mathsf{Pr}\left[\left|\mathbf{x}\mathbf{\Omega}^{\mathsf{T}}(\mathbf{A}^{\mathsf{T}}\mathbf{A})^{-1}\mathbf{\Omega}\mathbf{x}^{\mathsf{T}} - \mathbf{x}\mathbf{\Omega}^{\mathsf{T}}(\widetilde{\mathbf{A}}^{\mathsf{T}}\widetilde{\mathbf{A}})^{-1}\mathbf{\Omega}\mathbf{x}^{\mathsf{T}}\right| \le \varepsilon_0\right] \ge 1-\delta_0,$$

where $\varepsilon_0 = \frac{\varepsilon}{\sqrt{4r\ln(2/\delta)}}$ and $\delta_0 = \delta/2r$.

The first part of the proof follows simply as in [4]. More concretely, we have $\det(\mathbf{A}^\mathsf{T}\mathbf{A}) = \prod_i \sigma_i^2$, where $\sigma_1 \geq \cdots \geq \sigma_d \geq \sigma_{\mathsf{min}}$ are the singular values of $\mathbf{A}$. Let $\widetilde{\sigma}_1 \geq \cdots \geq \widetilde{\sigma}_d \geq \sigma_{\mathsf{min}}$ be its singular value for $\widetilde{\mathbf{A}}$. Since the singular values of $\mathbf{A} - \widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{A}} - \mathbf{A}$ are the same, $\sum_i (\sigma_i - \widetilde{\sigma}_i) \leq 1$ using Linskii's theorem. Therefore,

$$\sqrt{\prod_i \frac{\widetilde{\sigma}_i^2}{\sigma_i^2}} \leq \exp\left(\frac{\varepsilon}{32\sqrt{r\log(2/\delta)}\log(r/\delta)}\right)\sum_i(\widetilde{\sigma}_i - \sigma_i) \leq e^{\varepsilon_0/2}.$$

The second part of the proof is slightly more involved. Each row i of the published matrix is distributed identically and is constructed by multiplying an n-dimensional vector $\mathbf{\Omega}_i$ that has entries picked from a normal distribution $\mathcal{N}(0,1)$. Note that $\mathbb{E}[\mathbf{\Omega}_{i:}] = \mathbf{0}^n$ and $\mathsf{COV}(\mathbf{\Omega}_{i:}) = \mathbb{I}$. Then

$$\mathbf{x}\mathbf{\Omega}^\mathsf{T}(\mathbf{A}^\mathsf{T}\mathbf{A})^{-1}\mathbf{\Omega}\mathbf{x}^\mathsf{T} - \mathbf{x}\mathbf{\Omega}^\mathsf{T}(\widetilde{\mathbf{A}}^\mathsf{T}\widetilde{\mathbf{A}})^{-1}\mathbf{\Omega}\mathbf{x}^\mathsf{T} = \mathbf{x}\mathbf{\Omega}^\mathsf{T}\left[(\mathbf{A}^\mathsf{T}\mathbf{A})^{-1}(\mathbf{A}^\mathsf{T}\mathbf{E} + \mathbf{E}^\mathsf{T}\widetilde{\mathbf{A}})(\widetilde{\mathbf{A}}^\mathsf{T}\widetilde{\mathbf{A}})^{-1}\right]\mathbf{\Omega}\mathbf{x}^\mathsf{T}.$$

Using the singular value decomposition of $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\mathsf{T}$ and $\widetilde{\mathbf{A}} = \widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{U}}^\mathsf{T}$, this simplifies as

$$\left(\mathbf{x}\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\mathsf{T})\mathbf{e}_i\right)\left(\mathbf{v}^\mathsf{T}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}^{-2}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{x}^\mathsf{T}\right) + \left(\mathbf{x}\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-2}\mathbf{V}^\mathsf{T})\mathbf{v}\right)\left(\mathbf{e}_i^\mathsf{T}(\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{x}^\mathsf{T}\right).$$

Since $\mathbf{x} \sim \mathbf{A}^\mathsf{T}\mathbf{y}$, where $\mathbf{y} \sim \mathcal{N}(0,1)$, we can further simplify it as

$$\underbrace{\left(\mathbf{y}^\mathsf{T}\mathbf{A}\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\mathsf{T})\mathbf{e}_i\right)}_{t_1}\underbrace{\left(\mathbf{v}^\mathsf{T}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}^{-2}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{A}^\mathsf{T}\mathbf{y}\right)}_{t_2} + \underbrace{\left(\mathbf{y}^\mathsf{T}A\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-2}\mathbf{V}^\mathsf{T})\mathbf{v}\right)}_{t_3}\underbrace{\left(\mathbf{e}_i^\mathsf{T}(\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{A}^\mathsf{T}\mathbf{y}\right)}_{t_4}.$$

Now since $\|\widetilde{\mathbf{\Sigma}}\|_2, \|\mathbf{\Sigma}\|_2 \geq w$, plugging in the SVD of $\mathbf{A}$ and $\mathbf{A} - \mathbf{A}' = \mathbf{e}_i\mathbf{v}^\mathsf{T}$, and that every term t_i in the above expression is a linear combination of a Gaussian, i.e., each term is distributed as per $\mathcal{N}(0, \|t_i\|^2)$, we calculate $\|t_i\|$ as below.

$$\begin{aligned}
&\|(\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\mathsf{T})\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\mathsf{T})\mathbf{e}_i\|_2 \leq \|\mathbf{\Omega}^\mathsf{T}\|_2 \leq 1,\\
&\|(\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\mathsf{T})\mathbf{\Omega}^\mathsf{T}(\mathbf{V}\mathbf{\Sigma}^{-2}\mathbf{V}^\mathsf{T})\mathbf{v}\|_2 \leq \|\mathbf{\Omega}^\mathsf{T}\|_2\|\mathbf{\Sigma}^{-1}\|_2 \leq \frac{1}{\sigma_{\mathsf{min}}},\\
&\|v^\mathsf{T}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}^{-2}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{U}}^\mathsf{T} - \mathbf{v}\mathbf{e}_i^\mathsf{T})\|_2 \leq \|v^\mathsf{T}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}^{-2}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{U}}^\mathsf{T}\|_2 + \|v^\mathsf{T}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}^{-2}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{v}\mathbf{e}_i^\mathsf{T}\|_2\\
&\qquad \leq \frac{1}{\sigma_{\mathsf{min}}} + \frac{1}{\sigma_{\mathsf{min}}^2},\\
&\|e_i^\mathsf{T}(\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{U}}^\mathsf{T} - \mathbf{v}\mathbf{e}_i^\mathsf{T})\|_2 \leq \|\mathbf{e}_i^\mathsf{T}(\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}(\widetilde{\mathbf{V}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{U}}^\mathsf{T}\|_2 + \|\mathbf{e}_i^\mathsf{T}(\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^\mathsf{T})\mathbf{\Omega}\mathbf{v}\mathbf{e}_i^\mathsf{T}\|_2\\
&\qquad \leq 1 + \frac{1}{\sigma_{\mathsf{min}}},
\end{aligned}$$

where $\sigma_{\mathsf{min}} = \left(\frac{\sqrt{r\log(2/\delta)}\log(r/\delta)}{\varepsilon}\right)$. Using the concentration bound on the Gaussian distribution, each term, t_1, t_2, t_3, and t_4, is less than $\|t_i\|\ln(4/\delta_0)$ with probability $1-\delta_0/2$. From the fact that $2\left(\frac{1}{\sigma_{\mathsf{min}}} + \frac{1}{\sigma_{\mathsf{min}}^2}\right)\ln(4/\delta_0) \leq \varepsilon_0$, we have the second part of the proof.

We now prove that the second variant preserves privacy if the singular values of the streamed matrix follows the hypothesis of the theorem. We first compute the PDF of the published matrix.

We start by computing the probability density function when the underlying multivariate Gaussian distribution is $\mathcal{N}(0,\mathbb{I})$. The case for arbitrary positive definite covariance matrix follows like the transition from identity to arbitrary positive definite covariance matrices in the multivariate Gaussian distribution. Let $\mathbf{g}_1, \cdots, \mathbf{g}_r \sim \mathcal{N}(0,\mathbb{I})$ be r i.i.d multivariate Gaussian distribution, i.e., $\mathbf{g}_{ij} \sim \mathcal{N}(0,1)$ for

$1 \leq i \leq r, 1 \leq j \leq n$. The distribution we are interested in is $\mathbf{\Phi} = \sum_{i=1}^{r} \mathbf{g}_i \mathbf{g}_i^{\mathsf{T}}$. We use the notation $\mathsf{PDF}(\mathbf{\Phi}; \mathbb{I})$ to denote the probability density function of $\mathbf{\Phi}$ when each random variable is picked using a normal distribution, i.e., when the covariance matrix of the random variables is $\mathbb{I}$.

Using the chain rule, the joint distribution of the entries of $\mathbf{\Phi}$ is as follows.

$$\begin{aligned}\mathsf{PDF}(\mathbf{\Phi}; \mathbb{I}) = \mathsf{PDF}(\langle \mathbf{g}_1, \mathbf{g}_1 \rangle; \mathbb{I}) \mathsf{PDF}((\langle \mathbf{g}_2, \mathbf{g}_1 \rangle, \langle \mathbf{g}_2, \mathbf{g}_2 \rangle) | \langle \mathbf{g}_1, \mathbf{g}_1 \rangle; \mathbb{I}) \cdots \\ \mathsf{PDF}((\langle \mathbf{g}_r, \mathbf{g}_1 \rangle, \cdots \langle \mathbf{g}_r, \mathbf{g}_r \rangle) | \mathbf{\Phi}_{[r-1]}; \mathbb{I}). \end{aligned} \tag{2}$$

There are $r(r+1)/2$ distinct entries, $\langle \mathbf{g}_1, \mathbf{g}_1 \rangle, (\langle \mathbf{g}_2, \mathbf{g}_1 \rangle, \langle \mathbf{g}_2, \mathbf{g}_2 \rangle), \cdots, (\langle \mathbf{g}_r, \mathbf{g}_1 \rangle, \cdots \langle \mathbf{g}_r, \mathbf{g}_r \rangle)$. Our aim is to compute each individual term in the product form of the above chain rule. For this, we first analyze and understand the properties of the distribution of $\mathbf{h}_{i-1}^{\mathsf{T}} = (\langle \mathbf{g}_i, \mathbf{g}_1 \rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{i-1} \rangle)$. Then we use the fact that there is a transformation of Jacobian one from $\left(\langle \mathbf{g}_i, \mathbf{g}_1 \rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{(i-1)} \rangle, \langle \mathbf{g}_i, \mathbf{g}_i \rangle - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1} \right)$ to $\left(\langle \mathbf{g}_i, \mathbf{g}_1 \rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{(i)} \rangle \right)$ to compute each term in the chain rule (see, [32] and [33]).

Distribution of $\mathbf{h}_{i-1}$. We first prove that $\mathbf{h}_{i-1}$ is an $(i-1)$-variate Gaussian distribution. Since the covariance matrix is $\mathbb{I}$, and $\mathbf{g}_{11}, \cdots, \mathbf{g}_{1n}, \cdots, \mathbf{g}_{r1}, \cdots, \mathbf{g}_{rn}$ are i.i.d. $\mathcal{N}(0,1)$, from the elementary property of linear functions of normal variables, conditional on $\mathbf{g}_{kj}$ for $1 \leq k \leq i-1$ and $1 \leq j \leq n$, $\mathbf{h}_{i-1}$ is $(i-1)$-variate Gaussian distribution with

$$\mathbf{\Phi}_{[i]} = \begin{pmatrix} \langle \mathbf{g}_1, \mathbf{g}_1 \rangle & \cdots & \langle \mathbf{g}_1, \mathbf{g}_i \rangle \\ \vdots & \ddots & \vdots \\ \langle \mathbf{g}_i, \mathbf{g}_1 \rangle & \cdots & \langle \mathbf{g}_i, \mathbf{g}_i \rangle \end{pmatrix}$$

Now $\mathbf{g}_{11}, \cdots, \mathbf{g}_{rn}$, for every $j = 1, \cdots, n$ are mutually independent; therefore, we have the covariance

$$\mathsf{COV}(\mathbf{h}_{i-1}, \mathbf{g}_{ij}) = \left[\mathsf{COV}\left(\langle \mathbf{g}_i, \mathbf{g}_1 \rangle \mathbf{g}_{ij} \right), \cdots, \mathsf{COV}\left(\langle \mathbf{g}_i, \mathbf{g}_{(i-1)} \rangle \mathbf{g}_{ij} \right) \right] = (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j})^{\mathsf{T}}$$

and $\mathbb{E}[|\mathbf{h}_{i-1}\rangle\langle\mathbf{h}_{i-1}||\mathbf{g}_{kj}] = \mathbf{\Phi}_{[i-1]}$ for $1 \leq j < i$. This implies

$$\mathsf{COV}\left[\mathbf{h}_{i-1}, \mathbf{g}_{ij} - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} (\mathbf{g}_1 j, \cdots \mathbf{g}_{i-1,j})^{\mathsf{T}} | \mathbf{g}_{kj} \right] = 0 \qquad \forall 1 \leq k \leq i-1, \tag{3}$$

as the left hand side equals $(\mathbf{g}_1 j, \cdots \mathbf{g}_{i-1,j})^{\mathsf{T}} - \mathbf{\Phi}_{[i-1]} \mathbf{\Phi}_{[i-1]}^{-1} (\mathbf{g}_1 j, \cdots \mathbf{g}_{i-1,j})^{\mathsf{T}}$.

This implies that $\mathbf{h}_{i-1}$ is independent of $\sum_{j=1}^{k} \left(\mathbf{g}_{ij} - (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j})^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} (\mathbf{g}_{1j, \cdots, \mathbf{g}_{i-1,j}}) \right)^2$. [33] proved that

$$\sum_{j=1}^{k} \left(\mathbf{g}_{ij} - (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j})^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j}) \right)^2 \sim \chi^2_{n-i+1}, \tag{4}$$

the standard χ^2 distribution with $(n-i+1)$ degrees of freedom.

Computing every term in the chain rule. From the fact that $\mathbf{h}_{i-1}$ is a $(i-1)$-variate Gaussian distribution, equation (3), equation (4), and the identity

$$\Delta(\mathbf{\Phi}_{[i]}) = \Delta(\mathbf{\Phi}_{[i-1]}) \sum_{j=1}^{k} \left(\mathbf{g}_{ij} - (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j})^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} (\mathbf{g}_{1j}, \cdots, \mathbf{g}_{i-1,j}) \right)^2,$$

where $\Delta(\cdot)$ denotes the determinant, we first calculate the joint pdf of

$$\begin{aligned}
&\left(\langle \mathbf{g}_i, \mathbf{g}_1\rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{(i-1)}\rangle, \langle \mathbf{g}_i, \mathbf{g}_i\rangle - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1}\right)^{\mathsf{T}} \\
&\quad = \frac{\exp\left(-\frac{1}{2}\left(\langle \mathbf{g}_i, \mathbf{g}_1\rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{(i-1)}\rangle\right)^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \left(\langle \mathbf{g}_i, \mathbf{g}_1\rangle, \cdots, \langle \mathbf{g}_i \mathbf{g}_{(i-1)}\rangle\right)\right)}{(2\pi)^{(i-1)/2} \Delta(\mathbf{\Phi}_{[i-1]})^{1/2}} \\
&\qquad \times \frac{\exp\left(-\frac{\langle \mathbf{g}_i, \mathbf{g}_i\rangle - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1}}{2}\right) \left(\langle \mathbf{g}_i, \mathbf{g}_i\rangle - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1}\right)^{(n-i+1)/2-1}}{2^{(n-i+1)/2} \Gamma((n-i+1)/2)} \\
&\quad = \frac{\exp\left(-\frac{\mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1}}{2}\right) \Delta(\mathbf{\Phi}_i)^{(n-i-1)/2}}{2^{m/2} \pi^{(i-1)/2} \Gamma((n-i+1)/2) \Delta(\mathbf{\Phi}_{[i-1]})^{(n-i)/2}} \\
&\quad = \mathsf{PDF}((\langle \mathbf{g}_i, \mathbf{g}_1\rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{i-1}\rangle, \langle \mathbf{g}_i, \mathbf{g}_i\rangle) | \mathbf{\Phi}_{[i-1]}),
\end{aligned} \tag{5}$$

where the last step uses the fact that the transformation from $\left(\mathbf{h}_i^{\mathsf{T}}, \langle \mathbf{g}_i, \mathbf{g}_i\rangle - \mathbf{h}_{i-1}^{\mathsf{T}} \mathbf{\Phi}_{[i-1]}^{-1} \mathbf{h}_{i-1}\right)$ to $\left(\langle \mathbf{g}_i, \mathbf{g}_1\rangle, \cdots, \langle \mathbf{g}_i, \mathbf{g}_{(i)}\rangle\right)$ is one-to-one with Jacobian 1.

Computing the joint distribution of $\mathbf{\Phi}$. We are done except for putting in the values of every terms in the chain rule. A simple arithmetic followed by plugging equation (5) in equation (2) gives the closed formed expression of the pdf of $\mathbf{\Phi}$ as

$$\begin{aligned}
&\frac{\exp(-\mathrm{Tr}(\mathbf{\Phi})/2) \Delta(\mathbf{\Phi})^{(n-r-1)/2}}{2^{rn/2} \pi^{\sum_i (i-1)/2} \prod_{i=1}^r \Gamma((n-i+1)/2)} \times \prod_{i=1}^r \left(\frac{\Delta(\mathbf{\Phi}_{[i]})^{(n-i-1)/2}}{\Delta(\mathbf{\Phi}_{[i-1]})^{(n-i)/2}}\right) \\
&\qquad = \frac{\exp(-\mathrm{Tr}(\mathbf{\Phi})/2) \Delta(\mathbf{\Phi})^{(n-r-1)/2}}{2^{rn/2} \pi^{n(n-1)/4} \prod_{i=1}^r \Gamma((n-i+1)/2)}.
\end{aligned}$$

Let $\mathbf{\Phi}$ be the matrix formed in the manner as above with random vectors picked using the distribution that is defined by $\mathcal{N}(\mathbf{0}^n, \mathbf{\Sigma})$. Let $\mathbf{\Phi}' = \mathbf{\Phi}\mathbf{A}$. Then we can write the probability distribution function of $\mathbf{\Phi}'$ as follows

$$\begin{aligned}
\mathsf{PDF}_{\mathbf{\Phi}'}(\mathbf{\Phi}') &\propto \mathsf{PDF}_{\mathbf{\Phi}}(\mathbf{\Phi}'\mathbf{A}^{-1}) \\
&\propto \exp\left(\mathrm{Tr}\left(-\mathbf{\Sigma}^{-1}\mathbf{\Phi}\mathbf{A}^{-1}/2\right)\right) \\
&= \exp\left(\mathrm{Tr}\left(-\mathbf{A}^{-1}\mathbf{\Sigma}^{-1}\mathbf{\Phi}\right)/2\right),
\end{aligned}$$

where $\mathbf{A}^{-1}$ is the pseudo-inverse of $\mathbf{A}$. This is not distributed like the distribution of $\mathbf{\Phi}$.

We can now prove the privacy guarantee. Let $\delta_0 = \delta/r$. Let $\mathbf{A}$ and $\widetilde{\mathbf{A}}$ be a pair of neighbouring matrices that form the stream. From our definition of neighbouring matrices, $\mathbf{A} - \widetilde{\mathbf{A}} = \mathbf{E} = |\mathbf{v}\rangle\langle \mathbf{e}_i|$ for an unit vector $|\mathbf{v}\rangle$ and some i. The published matrices corresponding to the two neighbouring matrices have the following probability density function

$$\begin{aligned}
\mathsf{PDF}(\mathbf{\Phi}\mathbf{A}; \mathbb{I}) &= \frac{\exp(-\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{\Phi})/2) \Delta(\mathbf{\Phi})^{(n-r-1)/2}}{2^{rn/2} \pi^{n(n-1)/4} \Delta(\mathbf{A})^{r/2} \prod_{i=1}^r \Gamma((n-i+1)/2)} = C \frac{\exp(-\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{\Phi})/2)}{\Delta(\mathbf{A})^{r/2}}, \\
\mathsf{PDF}(\mathbf{\Phi}\widetilde{\mathbf{A}}; \mathbb{I}) &= \frac{\exp(-\mathrm{Tr}((\widetilde{\mathbf{A}}^{-1}\mathbf{\Phi})/2) \Delta(\mathbf{\Phi})^{(n-r-1)/2}}{2^{rn/2} \pi^{n(n-1)/4} \Delta(\widetilde{\mathbf{A}})^{r/2} \prod_{i=1}^r \Gamma((n-i+1)/2)} = C \frac{\exp(-\mathrm{Tr}((\widetilde{\mathbf{A}}^{\mathsf{T}}\widetilde{\mathbf{A}})^{-1}\mathbf{\Phi})/2)}{\Delta(\widetilde{\mathbf{A}})^{r/2}},
\end{aligned}$$

where $C = \Delta(\mathbf{\Phi})^{(n-r-1)/2} / \left(2^{rn/2}\pi^{n(n-1)/4}\prod_{i=1}^{r}\Gamma((n-i+1)/2)\right)$. It is straightforward to see that combination of the following proves differential privacy of the published matrix:

$$\exp(-\varepsilon/r) \leq \sqrt{\frac{\widetilde{\Delta}(\mathbf{A})}{\widetilde{\Delta}(\widetilde{\mathbf{A}})}} \leq \exp(\varepsilon/r) \quad \text{and} \quad \mathsf{Pr}\left[|\mathrm{Tr}\left(\left(\mathbf{A}^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right)| \leq \varepsilon\right] \geq 1-\delta. \tag{6}$$

Let $\sigma_1 \geq \cdots, \geq \sigma_d \geq \sigma_{\mathsf{min}}$ be the singular values of $\mathbf{A}$. Let $\widetilde{\sigma}_1, \geq \cdots, \geq \widetilde{\sigma}_d \geq \sigma_{\mathsf{min}}$ be the singular value for $\widetilde{\mathbf{A}}$. Since the singular values of $\mathbf{A} - \widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{A}} - \mathbf{A}$ are the same, $\sum_{i\in G}(\sigma_i - \widetilde{\sigma}_i) \leq 1$ using Linskii's theorem, where G is the set of indices for which $\sigma_i > \widetilde{\sigma}_i$. The first bound follows similarly as in Blocki *et al.* [4]. For the second bound required for the privacy, we first bound the following

$$\begin{aligned}\mathrm{Tr}\left(\left(\mathbf{A}^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) &= \mathrm{Tr}\left(\left(\mathbf{A}^{-1}(\widetilde{\mathbf{A}})(\widetilde{\mathbf{A}})^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) \\ &= \mathrm{Tr}\left(\left(\mathbf{A}^{-1}(\mathbf{A}+\mathbf{E})\widetilde{\mathbf{A}}^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) = \mathrm{Tr}\left(\left(\mathbf{A}^{-1}\mathbf{E}\widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right).\end{aligned}$$

Using the singular value decomposition of $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathsf{T}}$ and $\widetilde{\mathbf{A}} = \widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{V}}^{\mathsf{T}}$, and the fact that $\mathbf{E} = \mathbf{v}\mathbf{e}_i^{\mathsf{T}}$ for some i, we can further solve the above expression.

$$\begin{aligned}\mathrm{Tr}\left(\left(\mathbf{A}^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) &= \mathrm{Tr}\left(\left(\mathbf{A}^{-1}\mathbf{E}\widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) \\ &= \mathrm{Tr}\left(\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}\mathbf{v}\mathbf{e}_i^{\mathsf{T}}\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^{\mathsf{T}}\mathbf{\Phi}\right) \\ &= \sum_{j=1}^{r}\mathrm{Tr}\left(\mathbf{g}_j^{\mathsf{T}}\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}\mathbf{v}\mathbf{e}_i^{\mathsf{T}}\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^{\mathsf{T}}\mathbf{g}_j\right).\end{aligned}$$

Fix a j. We bound the following.

$$\left|\mathrm{Tr}\left(\mathbf{g}_j^{\mathsf{T}}\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}\mathbf{v}\mathbf{e}_i^{\mathsf{T}}\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^{\mathsf{T}}\mathbf{g}_j\right)\right|. \tag{7}$$

We now look at each term in the above expression. $\mathbf{g}_j^{\mathsf{T}}\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}\mathbf{e}_i$ as $\mathcal{N}(0, \|\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}\mathbf{e}_i\|^2)$, and $\mathbf{v}^{\mathsf{T}}\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^{\mathsf{T}}\mathbf{g}_j$ as $\mathcal{N}(0, \|\mathbf{v}^{\mathsf{T}}\widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}^{-1}\widetilde{\mathbf{V}}^{\mathsf{T}}\|^2)$. Since $\mathbf{v}$ and $\mathbf{e}_i$ are unit vectors, the norm of the above four quantities are less than $1/\sigma_{\mathsf{min}}$, and $1 + 1/\sigma_{\mathsf{min}}$, respectively.

Therefore, from the concentration inequality of Gaussian distribution, we have

$$\mathsf{Pr}\left[(7) \leq 2\left(\frac{1}{\sigma_{\mathsf{min}}} + \frac{1}{\sigma_{\mathsf{min}}^2}\right)\ln(4/\delta_0) \leq \varepsilon\right] \geq 1 - \delta_0.$$

Taking union bound, we have with probability $1-\delta$, $-\varepsilon \leq \mathrm{Tr}\left(\left(\mathbf{A}^{-1} - \widetilde{\mathbf{A}}^{-1}\right)\mathbf{\Phi}\right) \leq \varepsilon$.

□

Differences between PSG_1 and Blocki *et al.* [4]. At high level, PSG_1 has some resemblance to the mechanism of Blocki *et al.* [4] and Upadhyay [39] if we use random Gaussian matrix as $\mathbf{\Omega}$. However, the analogy ends here, for eg., Blocki *et al.* [4] and Upadhyay [39] perform an affine transformation to convert the private matrix into a set of $\{\sqrt{w/n}, 1\}^n$ vectors, while we perform perturbation to raise the singular values before invoking PSG (see Sections 4 and 5). The mechanism of [4, Algorithm 3] does not give a guarantee that the singular values of $\mathbf{A}^{\mathsf{T}}\mathbf{A}$ and their published matrix is close or their eigenvalues are comparable. In other words, it does not give a LRA. Apart from these major differences, there are couple of subtle differences: (i) they project the entries of the columns of private matrix to a higher dimensional space; here, we perform embedding to a lower dimensional subspace in the similar vein as other applications of dimensionality reduction, and (ii) their mechanism uses multiple passes over the input matrix (they require at least two-passes over the input matrix even with the streaming algorithms for computing the singular value decomposition (SVD). So, we cannot use their mechanism in any of the problems we study in this paper.

On input parameters $\alpha, \beta, \varepsilon, \delta$, the target rank k, set $w = \left(ck\varepsilon^{-1}\ln(k/\delta)\right)$ for a global constant c. Pick a $2n \times k$ standard Gaussian matrix $\mathbf{\Omega}$. On input an $n \times n$ matrix $\mathbf{A}$ of rank r, the mechanism does the following:

Range Finding. Compute $\mathbf{Y_A} = \begin{pmatrix} w\mathbb{I} & A \end{pmatrix} \mathbf{\Omega}$ by computing $\begin{pmatrix} w\mathbf{e}_i & \mathbf{A}_{i:} \end{pmatrix} \mathbf{\Omega}$ for all streams $\mathbf{A}_{i:}, 1 \leq i \leq n$ and appending them row-wise. Let $\Pi_{\mathbf{Y_A}} = \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}$ be the projection matrix corresponding to the range of $\mathbf{Y_A}$.

Projection: When the whole matrix is streamed, the curator does the following:

1. Let the (unknown) matrix $\mathbf{B} = \mathbf{\Psi}^\mathsf{T}\mathbf{A}_t\mathbf{\Psi}$. Use the minimal residual method to find a solution to $\mathbf{B}\mathbf{\Psi}_t^\mathsf{T}\mathbf{\Omega} = \mathbf{\Psi}_t^\mathsf{T}\mathbf{Y}_t$.
2. Compute the decomposition of $\mathbf{B}_t = \bar{\mathbf{U}}_t\mathbf{\Lambda}_t\bar{\mathbf{U}}_t^\mathsf{T}$, form the product $\hat{\mathbf{U}}_t = \mathbf{\Psi}_t\bar{\mathbf{U}}_t$, and publish $\widehat{\mathbf{U}}_t\mathbf{\Lambda}_t\widehat{\mathbf{U}}_t^\mathsf{T}$.

Figure 2: The Mechanism for k-rank Approximation

4 Low Rank Approximation

Blocki *et al.* [4] noted that their published matrix is neither close nor their eigenvalues are comparable to the private matrix. In other words, it does not give a LRA; therefore, we need a different approach. We use the prototype mentioned in Halko *et al.* [20], which was also used by Hardt and Roth [22] to improve the worst-case bound under a *low coherence* assumption. In this prototype, we construct a low-dimensional subspace that captures the action of the matrix (*range-finding*), and then restrict the matrix to that subspace to compute the required factorization (*projection*). More concretely, range finding finds a measurement matrix $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}$, where $\mathbf{\Omega}$ is a Gaussian matrix in our case and computes the orthonormal projection matrix $\Pi_{\mathbf{Y}}$ corresponding to the range defined by $\mathbf{Y}$; projection then computes a k-rank matrix $\mathbf{B} = \Pi_{\mathbf{Y}}\mathbf{A}$. From this exposition, it seems that privacy preserving algorithms are required for both the stages; however, we show that the two-step prototype can be replaced by a two-step algorithm in which the input matrix is explicitly needed only in the first step at the expense of privacy proof requiring both PSG_1 and PSG_2.

We first note that if $\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}$ is a LRA of $\mathbf{A}$, i.e., $\|\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^T\mathbf{A}\| \leq \eta$, then so is $\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}$. This is because $\|\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\| = \|\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A} + \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\| \leq \|\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}\| + \|\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A} - \mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\mathbf{A}\mathbf{\Psi}\mathbf{\Psi}^\mathsf{T}\| \leq 2\eta$. The crucial observation now for the single pass LRA when $\mathbf{\Omega}$ is a Gaussian matrix is that $\mathbf{\Omega}$, $\mathbf{Y}$, and the basis for the range of $\mathbf{Y}$ contains enough information to compute the matrix $\mathbf{B}$, i.e., we do not need $\mathbf{A}$ explicitly. The range-finding is private using Theorem 3.1(i), but as we reuse $\mathbf{\Omega}$, we have to rely on Theorem 3.1(ii) to prove the privacy of projection step. In order to simplify the presentation, we state our mechanism for symmetric matrices in Figure 2 (Hardt and Roth [23] and Kapralov and Talwar [28] also made this assumption) that computes a LRA of the SVD of $\mathbf{A}$.

The mechanism for LRA presented in Figure 2 assume that the matrix $\mathbf{A}$ is provided as the symmetric $\mathsf{rank}(\mathbf{A})$ matrix $\mathbf{A}'$ corresponding to $\mathbf{A}$, and stop updating the data structure once all the rows of $\mathbf{A}$ are streamed. This simplifies the presentation as well as the analysis. By the argument of [22, Fact 2.8], this leads to a depreciation of the privacy guarantee by half (both ε and δ). The analysis for error bound though has complications because the right and the left singular vectors of the original matrix might be different. Keeping this in mind, we present its analysis in the most general form.

Theorem 4.1. Let $\sigma_1 \geq \cdots \geq \sigma_{\mathsf{rk}(\mathbf{A})}$ be the singular values of $\mathbf{A}$. Then for an over-sampling parameter p with the most common choice being $p = k + 1$, there is a single-pass mechanism that computes k-rank

approximation $\bar{\mathbf{A}}$ using $O(k(n+d)\alpha^{-1}\kappa)$ bits while preserving (ε,δ)-differential privacy such that

$$(i)\quad \|\mathbf{A}-\bar{\mathbf{A}}\|_F \leq \left(1+\frac{k}{p-1}\right)^{1/2} \min_{rk(A')<k} \|\mathbf{A}-\mathbf{A}'\|_F + \frac{2k}{\varepsilon}\sqrt{\frac{(n+d)\ln(k/\delta)}{p}}, \qquad \text{and}$$

$$(ii)\quad \|\mathbf{A}-\bar{\mathbf{A}}\|_2 \leq \left(1+\frac{k}{p-1}\right)^{1/2} \sigma_{k+1} + \frac{e\sqrt{(k+p)\sum_{j>k}\sigma_j^2}}{p} + \frac{2\sqrt{k(n+d)\ln(k/\delta)}}{\varepsilon}.$$

The most common choice of p is $k+1$, which is what we use in Table 1.

Proof. The space complexity is easy to follow from our convention of bit complexity and because we need to store the matrix $\mathbf{\Omega}$ and the sketch. The privacy guarantee follows from Theorem 3.1 and noting that all the singular-values of the matrix on which $\mathbf{\Omega}$ is operated from the right is greater than the threshold required for the statement of the Theorem 3.1, and the distribution of $\mathbf{\Psi}^\mathsf{T}\mathbf{\Omega}$ is the same as that of the second variant as we reuse $\mathbf{\Omega}$ (this follows from [7]). Now, it follows from the proof of the second variant (Theorem 3.1) that it does not matter if we multiply $\mathbf{A}$ (or $\mathbf{A}^\mathsf{T}$) from left (or right, respectively) of vectors $\mathbf{g}_1,\cdots,\mathbf{g}_r$, i.e., $\sum_{i=1}^r \mathbf{A}|\mathbf{g}_i\rangle\langle\mathbf{g}_i| = \mathbf{A}\mathbf{\Phi}$ and $\sum_{i=1}^r |\mathbf{g}_i\rangle\langle\mathbf{g}_i|\mathbf{A}^\mathsf{T} = \mathbf{\Phi}\mathbf{A}^\mathsf{T}$ have the same distribution. Combining all these arguments, we have the distribution of the second step of projection stage is identical to the second variant in Figure 1, modulo some deterministic computation. Since, any arbitrary post-processing preserves differential privacy, we can now complete the proof by invoking Theorem 3.1. The privacy guarantee due to Theorem 3.1 requires the minimum singular value to be greater than $\frac{4\sqrt{k\log(2/\delta)}\log(k/\delta)}{\varepsilon}$ for the first variant and $\frac{k\log(k/\delta)}{\varepsilon}$ for the second variant. By our choice of w, the singular values of the streamed matrix to the algorithm for PSG are at least the eigenvalues of $\sqrt{w^2\mathbb{I}+\mathbf{A}^\mathsf{T}A}$, which are all greater than $\frac{16k\ln(k/\delta)}{\varepsilon}$. Since $\frac{4\sqrt{k\log(2/\delta)}\log(k/\delta)}{\varepsilon} \ll 16\frac{k\log(k/\delta)}{\varepsilon}$, the privacy guarantee follows from Theorem 3.1.

In more detail, it follows from Bura and Pfeiffer [7][1] and the proof in Theorem 3.1 that, for $\mathbf{\Phi} = \sum_{i=1}^r |\mathbf{g}_i\rangle\langle\mathbf{g}_i|$ and large enough n, the distribution of $\mathbf{\Psi}^\mathsf{T}\mathbf{\Omega}$ is

$$\frac{\exp(-\mathrm{Tr}(\mathbf{A}^{-1}(\mathbf{D}^{-1}\mathbf{R})(\mathbf{D}^{-1}\mathbf{R})^\mathsf{T}\mathbf{\Phi})/2)\Delta(\mathbf{\Phi}(\mathbf{D}^{-1}\mathbf{R})(\mathbf{D}^{-1}\mathbf{R})^\mathsf{T})^{(n-r-1)/2}}{2^{rn/2}\pi^{n(n-1)/4}\Delta(\mathbf{A})^{r/2}\prod_{i=1}^r\Gamma((n-i+1)/2)}.$$

The proof for the first expression of equation (6) is as before. Following the rest of the steps of Theorem 3.1, using Lemma 2.2, we compute for fixed $j\in[k+p]$, instead of equation (7), the second expression in equation (6) is bounded by the following expression

$$\left|\mathrm{Tr}\left(\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\mathsf{T}\mathbf{e}_i\mathbf{v}^\mathsf{T}\tilde{\mathbf{V}}\mathbf{\Lambda}^{-2}\tilde{\mathbf{V}}^\mathsf{T}\alpha_j\alpha_j^\mathsf{T}\right)\mathrm{Tr}\left((\mathbf{D}^{-1}\mathbf{R})(\mathbf{D}^{-1}\mathbf{R})^\mathsf{T}\right)\right|.$$

Since $\mathbf{R}$ is an orthonormal matrix, using the fact that $\mathrm{Tr}((\mathbf{D}^\mathsf{T}\mathbf{D})^{-1}) = \mathrm{Tr}((\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega})^{-1})$, Lemma 2.5 with $p=k+1$, and following the remaining steps of Theorem 3.1, we have the privacy result for computing $\mathbf{B}$. The complete proof follows from the discussion stated at the start of this proof.

Remark 1. We take the liberty to diverge a little to understand the intuition behind Bura and Pfeiffer [7]. One may skip this part without effecting the readability of the rest of this section. The result of Bura and Pfeiffer [7] uses advance statistical tools, but the intuition can be argued using some basic statistics. On the other hand, it is well known that the singular values of random matrices are notoriously hard to compute (see [34]). The basic reasoning behind their proof is the following line of argument. Since the entries of

[1] We use the following result of Bura and Pfeiffer [7]: for a random $n\times k$ normal matrix $\mathbf{\Omega}$, for large enough n, the vector formed by the entries of its left singular matrix is distributed normally with covariance matrix $(\mathbf{D}^{-1}\mathbf{R}\otimes\mathbb{I})(\mathbf{R}\mathbf{D}^{-1}\otimes\mathbb{I})^\mathsf{T}$, where $\mathbf{\Omega}=\mathbf{L}\mathbf{D}\mathbf{R}^\mathsf{T}$.

an $n \times r$ matrix $\mathbf{\Omega}$ is $\mathcal{N}(0,1)$, then for any orthogonal matrices $\mathbf{G} \in \mathbb{R}^{n\times n}$ and $\mathbf{R} \in \mathbb{R}^{r\times r}$, the entries of $\mathbf{G\Omega R}^\mathsf{T}$ is also i.d.d. normal. This can be seen by translating to the vector form of the matrix, i.e., $\mathbf{v} = \mathsf{vec}(\mathbf{V})$ represent rn vector with entries $\mathbf{v}_{i+nj} = \mathbf{V}_{ij}$. Then $\mathsf{vec}(\mathbf{G\Omega R}^\mathsf{T}) = (\mathbf{G} \otimes \mathbf{R})\mathsf{vec}(\mathbf{\Omega})$. Now $\mathbf{G} \otimes \mathbf{R}$ is also an orthogonal matrix, and multivariate Gaussian distribution is preserved if one multiply by an orthogonal matrix. Therefore, the distribution of the left singular vectors of $\mathbf{\Omega}$ is the same as $\mathbf{G\Omega R}^\mathsf{T}$. Consequently, for large enough n, the distribution of each singular vector is also spherically distributed. □

UTILITY GUARANTEE. In Section 4, we showed that the mechanism for symmetric matrices does what [20, Section 1.2] prototype algorithm achieves. To construct a mechanism for non-symmetric matrices, we construct two sketches $\mathbf{Y}_t$ and $\bar{\mathbf{Y}}_t$ corresponding to $\mathbf{A}_1$ and $\mathbf{A}_2$ using a single pass over $\mathbf{A}$ and using two Gaussian matrices $\mathbf{\Omega}$ and $\bar{\mathbf{\Omega}}$ of appropriate dimension, where $\mathbf{A}_1 = \begin{pmatrix} w\mathbb{I} & A \end{pmatrix}$ and $\mathbf{A}_2 = \begin{pmatrix} \mathbf{A}^\mathsf{T} & w\mathbb{I} \end{pmatrix}$ for appropriate dimension identity matrices in both $\mathbf{A}_1$ and $\mathbf{A}_2$. Basically, we do the following using a single pass over the matrix $\mathbf{A}$:

$$\mathbf{Y} := \begin{pmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{pmatrix} = \begin{pmatrix} w\mathbb{I}_n & \mathbf{A} \\ \mathbf{A}^\mathsf{T} & w\mathbb{I}_d \end{pmatrix} \begin{pmatrix} \mathbf{\Omega}_1 \\ \mathbf{\Omega}_2 \end{pmatrix},$$

where $\mathbb{I}_n$ is an $n \times n$ identity matrix.

Since $\begin{pmatrix} \mathbf{Y}_1 & \mathbf{Y}_2 \end{pmatrix}^\mathsf{T}$ corresponds to a symmetric matrix, we can use the steps used in the projection stage in Figure 2. Note that Clarkson and Woodruff [10] compute $\mathbf{A}^\mathsf{T}A$ in a single-pass over the matrix $\mathbf{A}$.

Keeping the most general case in mind, we first show that the left singular vectors have hardly any role to play in bounding the perturbation. We assume that we perform SVD. Let the SVD of $\mathbf{A}$ be $\mathbf{U\Sigma V}^\mathsf{T}$. In the following discussion, we compute the approximation of $\begin{pmatrix} w\mathbb{I} & \mathbf{A} \end{pmatrix}$ and denote it by $\mathbf{A}$. This is because $\begin{pmatrix} w\mathbb{I} & \mathbf{A} \end{pmatrix}$ is more manageable and any upper bound on the approximation of this matrix is an upper bound on the approximation of the original matrix. The actual bound as computed in Figure 2 can be obtained by simply performing the computation on the singular value decomposition as performed in the last step of mechanism and using the sub-additivity of norms. From the discussion above, we know that $\bar{\mathbf{A}} = \Pi_\mathbf{Y}\mathbf{A}$; therefore, we need to bound $\|(\mathbb{I} - \Pi_\mathbf{Y})\mathbf{A}\|$, where, unless specified, in this section $\|\cdot\|$ refers to both the Frobenius as well as the spectral norm. From the Hölder's inequality on the second moment, we have

$$\mathbb{E}[\|(\mathbb{I} - \Pi_\mathbf{Y})A\|_F] \leq \left(\mathbb{E}\left[\|(\mathbb{I} - \Pi_\mathbf{Y})A\|_F^2\right]\right)^{1/2}. \tag{8}$$

We now bound $\| (\mathbb{I} - \Pi_\mathbf{Y})A \|$. Let $\mathbf{\Lambda} = \sqrt{\mathbf{\Sigma}^2 + w^2\mathbb{I}}$. Let $\mathbf{\Lambda}_1$ denote the diagonal matrix formed by the first k singular values and $\mathbf{\Lambda}_2$ be the diagonal matrix for the rest of the singular values. We decompose $\mathbf{V}^\mathsf{T}$ similarly. Let the matrix formed by the first k rows of $\mathbf{V}^\mathsf{T}$ be $\mathbf{V}_1^\mathsf{T}$ and by the rest of the rows be $\mathbf{V}_2^\mathsf{T}$.

Left singular vectors have essentially no role in the approximation bound. Recall that $\mathbf{\Omega}$ is an $2n \times k$ matrix; therefore, $\mathbf{Y} = \mathbf{A\Omega} = \mathbf{U}\begin{pmatrix} \mathbf{\Lambda}_1\mathbf{V}_1^\mathsf{T}\mathbf{\Omega} & \mathbf{\Lambda}_2\mathbf{V}_2^\mathsf{T}\mathbf{\Omega} \end{pmatrix}^\mathsf{T}$. It would be useful to consider the first k rows of $\mathbf{Y}$ to be the one that mimics the action of $\mathbf{A}$ and the rest of the rows of $\mathbf{Y}$ as a small perturbation that we wish to bound. We first prove that the left singular vectors have essentially no role to play in bounding the error. Let $\mathbf{A}' = \mathbf{UA}$, then the following chain of equalities are straightforward.

$$\|(\mathbb{I} - \Pi_\mathbf{Y})\mathbf{A}\| = \|\mathbf{U}^\mathsf{T}(\mathbb{I} - \Pi_\mathbf{Y})\mathbf{A}\| = \|\mathbf{U}^\mathsf{T}(\mathbb{I} - \Pi_\mathbf{Y})\mathbf{UA}'\| = \|\mathbb{I}\mathbf{A}' - \mathbf{U}^\mathsf{T}\Pi_\mathbf{Y}\mathbf{UA}'\|. \tag{9}$$

Now note that the projection matrix corresponding to a matrix $\mathbf{Y}$ is uniquely defined by $\mathsf{range}(\mathbf{Y})$, the range of $\mathbf{Y}$. Therefore, $\mathsf{range}(\mathbf{U}^\mathsf{T}\Pi_\mathbf{Y}\mathbf{U}) = \mathbf{U}^\mathsf{T}\mathsf{range}(\Pi_\mathbf{Y}) = \mathsf{range}(\mathbf{U}^\mathsf{T}\Pi_\mathbf{Y})$.

Therefore, $\|\mathbf{A}' - \mathbf{U}^\mathsf{T}\Pi_\mathbf{Y}\mathbf{UA}'\| = \|(\mathbb{I} - \Pi_{\mathbf{U}^\intercal\mathbf{Y}})\mathbf{A}'\|$. A useful way to understand the above expression is to view this geometrically and recall that unitary are just rotation in the space: projection by an unitary, followed by any projection operator, followed by the inverse of unitary brings us to the same space as projection by an operator followed by the inverse of the unitary.

Finding and bounding an appropriate perturbed matrix. We now use the identity that, for two operators $\mathbf{O}_1$ and $\mathbf{O}_2$, if the range of $\mathbf{O}_1$ is a subset of the range of $\mathbf{O}_2$, then the projection of any matrix using $\mathbf{O}_1$ will have all its norm smaller than the projection by $\mathbf{O}_2$. More concretely, we find a matrix $\mathbf{C}$ such that its range is a strict subset of the range of $\mathbf{U}^\mathsf{T}\mathbf{Y}$. We obtain this matrix by flattening out the first k rows of $\mathbf{U}^\mathsf{T}\mathbf{Y}$. This is in correspondence with our earlier observation that the first k rows mimic the action of $\mathbf{A}$ and other rows are the perturbation that we wish to bound. Since the first k rows of $\mathbf{U}^\mathsf{T}\mathbf{Y}$ is $\mathbf{\Lambda}_1\mathbf{V}_1^\mathsf{T}\mathbf{\Omega}$, let $\mathbf{C} := \mathbf{U}^\mathsf{T}\mathbf{Y}\mathbf{\Omega}^{-1}\mathbf{V}_1\mathbf{\Lambda}_1^{-1}$. The rows corresponding to the perturbation are $\mathbf{\Lambda}_2\mathbf{V}_2\mathbf{\Omega}$. Thus, $\mathbf{C} = \begin{pmatrix}\mathbb{I} & \mathbf{\Lambda}_2\mathbf{V}_2^\mathsf{T}\mathbf{V}_1\mathbf{\Lambda}_1^{-1}\end{pmatrix}^\mathsf{T}$.

Let us denote by $\mathbf{S} = \mathbf{\Lambda}_2\mathbf{V}_2^\mathsf{T}\mathbf{V}_1\mathbf{\Lambda}_1^{-1}$. It is not difficult to see that $\mathsf{range}(\mathbf{C}) \subset \mathsf{range}(\mathbf{U}^\mathsf{T}\mathbf{Y})$. Moreover, $\Pi_\mathbf{C} \preceq \mathbb{I}$, $\Pi_{\mathbf{U}^\intercal\mathbf{Y}}\Pi_\mathbf{C}\Pi_{\mathbf{U}^\intercal\mathbf{Y}} \preceq \Pi_{\mathbf{U}^\intercal\mathbf{Y}}$. This follows from the fact that $\mathsf{range}(\mathbf{C}) \subset \mathsf{range}(\mathbf{U}^\mathsf{T}\mathbf{Y})$ and the following derivation

$$\Pi_{\mathbf{U}^\intercal\mathbf{Y}} \succeq \Pi_{\mathbf{U}^\intercal\mathbf{Y}}\Pi_\mathbf{C}\Pi_{\mathbf{U}^\intercal\mathbf{Y}} = \Pi_\mathbf{C}\Pi_{\mathbf{U}^\intercal\mathbf{Y}} = (\Pi_{\mathbf{U}^\intercal\mathbf{Y}}\Pi_\mathbf{C})^\mathsf{T} = \Pi_\mathbf{C}.$$

An immediate result of the above is the following:

$$\|(\mathbb{I} - \Pi_{\mathbf{U}^\intercal\mathbf{Y}})\mathbf{A}'\| \leq \|(\mathbb{I} - \Pi_\mathbf{C})\mathbf{A}'\|. \tag{10}$$

Since, $\Pi_\mathbf{C} = \mathbf{C}(\mathbf{C}^\mathsf{T}\mathbf{C})^{-1}\mathbf{C}^\mathsf{T}$, we have the following set of derivations.

$$\begin{aligned}
\Pi_\mathbf{C} &= \begin{pmatrix}\mathbb{I}\\ \mathbf{S}\end{pmatrix}\left[\begin{pmatrix}\mathbb{I} & \mathbf{S}^\mathsf{T}\end{pmatrix}\begin{pmatrix}\mathbb{I}\\ \mathbf{S}\end{pmatrix}\right]^{-1}\begin{pmatrix}\mathbb{I} & \mathbf{S}^\mathsf{T}\end{pmatrix} = \begin{pmatrix}\mathbb{I}\\ \mathbf{S}\end{pmatrix}\left[(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})\right]^{-1}\begin{pmatrix}\mathbb{I} & \mathbf{S}^\mathsf{T}\end{pmatrix} \\
&= \begin{pmatrix}\mathbb{I}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\\ \mathbf{S}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\end{pmatrix}\begin{pmatrix}\mathbb{I} & \mathbf{S}^\mathsf{T}\end{pmatrix} = \begin{pmatrix}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1} & (\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T}\\ \mathbf{S}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1} & \mathbf{S}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T}\end{pmatrix} \\
&\succeq \begin{pmatrix}(\mathbb{I} - \mathbf{S}^\mathsf{T}\mathbf{S}) & (\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T}\\ \mathbf{S}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1} & 0\end{pmatrix},
\end{aligned}$$

where the last inequality uses the fact that $\mathbb{I} - \mathbf{S}^\mathsf{T}\mathbf{S} \preceq (\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}$ and $\mathbf{S}(\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T} \succeq 0$. Therefore,

$$\mathbb{I} - \Pi_\mathbf{C} \preceq \begin{pmatrix}\mathbf{S}^\mathsf{T}\mathbf{S} & \mathbb{I} - (\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T}\\ \mathbb{I} - ((\mathbb{I} + \mathbf{S}^\mathsf{T}\mathbf{S})^{-1}\mathbf{S}^\mathsf{T})^\mathsf{T} & \mathbb{I}\end{pmatrix}.$$

Conjugating $\mathbb{I} - \Pi_\mathbf{C}$ with $\mathbf{\Lambda}$, and applying the fact that for every positive definite matrix, $P = \begin{pmatrix}\mathbf{X} & \mathbf{Y}\\ \mathbf{Y}^\mathsf{T} & \mathbf{Z}\end{pmatrix}$, we have $\|P\| \leq \|\mathbf{X}\| + \|\mathbf{Z}\|$, we get

$$\|(\mathbb{I} - \Pi_\mathbf{C})\mathbf{A}'\| \leq \|\mathbf{S}^\mathsf{T}\mathbf{S}\mathbf{A}'\| + \|\mathbf{A}'\| \tag{11}$$

for any norm. From here on, it is easy arithmetic to show that

$$\|(\mathbb{I} - \Pi_\mathbf{C})\mathbf{A}'\| \leq \sqrt{\|\mathbf{\Lambda}_2'\| + \|\mathbf{\Lambda}_2'\mathbf{V}_2^\mathsf{T}\mathbf{\Omega}(\mathbf{V}_1^\mathsf{T}\mathbf{\Omega})^{-1}\|} \tag{12}$$

for both the required norms. Using equation (9), equation (10) and equation (11), this gives us a bound on the approximation of matrix $\mathbf{A}'$. Till this point, our analysis closely follows the ideas of [20], accommodating the steps of our algorithm. Now, all that remains is to bound $\|\mathbf{\Lambda}_2\mathbf{V}_2^\mathsf{T}\mathbf{\Omega}(\mathbf{V}_1^\mathsf{T}\mathbf{\Omega})^{-1}\|$, and for this, we have to analyze the matrix $\mathbf{\Omega}$.

ERROR BOUND FOR FROBENIUS NORM We now exploit the rotational invariance of a Gaussian distribution. An important point to note is that $(\mathbf{\Omega}\mathbf{\Omega}^{\mathsf{T}})^{-1}$ exists and has a well defined trace. The first part of the right hand side of equation (12) is immediate. Thus, if we bound $\mathbb{E}[\|\mathbf{\Lambda}_2'\mathbb{V}_2^{\mathsf{T}}\mathbf{\Omega}(\mathbf{V}_1\mathbf{\Omega})^{-1}\|]$, we are done. This could be accomplished as below.

$$\begin{aligned}
\mathbb{E}[\|\mathbf{\Lambda}_2\mathbb{V}_2^{\mathsf{T}}\mathbf{\Omega}(\mathbf{V}_1\mathbf{\Omega})^{-1}\|] &\leq \sqrt{\mathbb{E}\left[\sum_{ij}|(\mathbf{\Lambda}_2)'_{ij}\Pi_{ij}(\mathbf{V}_1\mathbf{\Omega}^{-1})_{jj}|\right]} \\
&\leq \sqrt{\|\mathbf{\Lambda}_2'\|_F\|\mathbf{\Omega}^{-1}\|_F} \\
&= \sqrt{\|\mathbf{\Lambda}_2'\|_F\mathsf{Tr}\left((\mathbf{\Omega}\mathbf{\Omega}^{-1})^{\mathsf{T}}\right)} = \sqrt{\|\mathbf{\Lambda}_2'\|_F\mathsf{Tr}(\mathbf{\Omega}\mathbf{\Omega}^{\mathsf{T}})^{-1}} \\
&\leq \sqrt{\mathsf{Tr}(\mathbf{\Omega}\mathbf{\Omega}^{\mathsf{T}})^{-1}}\min_{rk(\mathbf{A}')\leq k}\|\mathbf{A}-\mathbf{A}'\|_F + \sqrt{(n+d)w\mathsf{Tr}(\mathbf{\Omega}\mathbf{\Omega}^{\mathsf{T}})^{-1}}.
\end{aligned}$$

The utility guarantee follows by plugging this value in equation (12), and combining equation (8) and the fact that $(\mathbf{\Omega}\mathbf{\Omega}^{\mathsf{T}})^{-1}$ has a well defined trace $k/(p-1)$ [32].

ERROR BOUND FOR SPECTRAL NORM In order to bound the second term, we use few well known facts in the theory of random matrices to simplify equation (12). In particular, using Lemma 2.3 and 2.4, and Holder's inequality, the statement of the theorem for the spectral norm follows. The utility bound then follows using the same arithmetic of representing $\mathbf{\Lambda}'$ in terms of $\mathbf{\Lambda}$ as done in the case of Frobenius norm. In more details, we first bound

$$\begin{aligned}
\mathbb{E}\left[\|\mathbf{\Lambda}_2'\mathbf{V}_2^{\mathsf{T}}\mathbf{\Omega}(\mathbf{V}_1^{\mathsf{T}}\mathbf{\Omega})^{-1}\|\right] &\leq \|\mathbf{\Lambda}_2'\|\left(\mathbb{E}\left[\|(\mathbf{V}_1^{\mathsf{T}}\mathbf{\Omega})^{-1}\|_F^2\|\right]\right)^{1/2} + \|\mathbf{\Lambda}_2'\|_F\left(\mathbb{E}\left[\|(\mathbf{V}_1^{\mathsf{T}}\mathbf{\Omega})^{-1}\|\right]\right) \\
&\leq \|\mathbf{\Lambda}_2'\|\left(\mathbb{E}\left[\|\mathbf{\Omega}^{-1}\|_F^2\|\right]\right)^{1/2} + \|\mathbf{\Lambda}_2'\|_F\left(\mathbb{E}\left[\|\mathbf{\Omega}^{-1}\|\right]\right),
\end{aligned}$$

and then invoke Lemma 2.4 followed by the sub-additivity of norms. Making these substitution and on simplification, we get the bound stated in Theorem 4.1(ii).

$$(12) \leq \left(1+\frac{k}{p-1}\right)^{1/2}\|\mathbf{\Lambda}_2\|_2 + \frac{e\sqrt{(k+p)\sum_{j>k}\lambda_j^2}}{p} + \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}.$$

4.0.1 Tightness of Bounds and Comparison with Earlier Works

We compare our results with the best possible results in the non-private setting. Eckart and Young [18] have shown that the quantity $\min_{rk(\mathbf{A}')<k}\|\mathbf{A}-\mathbf{A}'\|_F$ in the first term of Theorem 4.1(i) is optimal. Likewise, Mirsky [29] proved that λ_{k+1} is the minimum spectral error for k-rank approximation. The second term in Theorem 4.1(ii) shows that we also pay for the Frobenius norm error when doing a unified analysis. However, when the oversampling parameter $p \approx k$, then the factor on λ_{k+1} is constant and that on the second term is of order $k^{-1/2}$. In fact, on closer analysis,

$$\|\mathbf{A}-\bar{\mathbf{A}}\|_2 \leq \left(1+\frac{k}{p-1}+\frac{e(\sqrt{(k+p)\min\{d,n\}}-k)}{p}\right)\lambda_{k+1} + \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon},$$

therefore, the error always lies within some polynomial factor of λ_{k+1}, modulo some additive error. As pointed out of Halko *et al.* [20], one can improve this by power-iteration, the method used by Hardt and Price [21] and Hardt and Roth [23] and multiple pass mechanism. However, it seems unlikely to improve it in a single-pass.

Kapralov and Talwar [28] showed a lower bound on additive error when $\delta = 0$ for neighbouring data differing by unit spectral norm. Our privacy proof depends strongly on the fact that $\delta \neq 0$. In fact, our bound is vacuous if $\delta = 0$. Though incomparable due to difference in the notion of neighbouring data, this separation gap further strengthen the belief that better bounds are possible for $\delta \neq 0$. Recently, Dwork *et al.* [17] also showed a bound in the online learning model, which is a different model of computation.

We compare our results with the works stated in Table 1 in more detail.

Chaudhary *et al.* [9]: They give low-rank approximation in the spectral norm. Additionally, they achieve $(\varepsilon, 0)$-differential privacy, which is only achieved by Kapralov and Talwar [28]. Their definition of neighbouring data sets can be (arguably) considered the most general in the sense that they consider two data sets neighbouring if they differ by at most one in the spectral norm. They use exponential distribution to sample a singular vector and give a heuristic, but practical implementation using Markov chain Monte-Carlo. On the negative side, their mechanism uses k rounds; therefore, it cannot be implemented in a streaming fashion. Since the notion of neighbouring data-sets and privacy guarantee achieved is different from that of ours, we believe our result is incomparable to that of Chaudhary *et al.* [9]. However, if we just concentrate on the additive error bound, they achieve a bound of order $O(nk/\varepsilon)$ compared to our bound $\tau \leq \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}$ (Theorem 4.1(ii)).

Hardt and Roth [22]: The authors use two passes over the input matrix; therefore, it does not fall in our one-pass streaming model of computation. They use the same notion of neighboring data-sets as we do in this paper. This makes their coherence conditions and notion of neighbouring data sets rotationally invariant. As argued by Blocki *et al.* [4], we achieve a better utility bound in the range finding step. Intuitively, this could be seen as a consequence of the absence of additive Gaussian noise. More concretely, Hardt and Roth [22] achieved an error bound of $\sqrt{kn}\log(k/\delta)/\varepsilon + \sqrt{\mu\|\mathbf{A}\|_F(n/d)^{1/2}\log(k/\delta)/\varepsilon}$. Their error bound depends on $\|\mathbf{A}\|_F$, which can be as large as $\sqrt{nd}$ for binary matrices in the worst case when the matrix is not as well-behaved as captured by low-coherence assumption. On the other hand, we achieve a bound that is independent of $\|\mathbf{A}\|_F$.

Hardt and Roth [23]: In some sense, this paper is based on Krylov subspace iteration combined with powering method of Halko *et al.* [20]. They define two data-sets as neighbouring in the same manner as in Hardt and Roth [23]. The coherence definition used in this paper depends on the maximum value of the left or right singular vectors, and is, therefore, rotationally variant. This work assumes that the singular value are well separated, i.e., the first and the k-th singular value has a non-trivial separation, and give LRA in spectral norm. Their bound, however, depends on the rank of the input matrix. Their mechanism uses k rounds of subspace generation, each of which depends on the spectrum of the matrix and uses the power-iteration method [20]. This helps them in achieving better multiplicative bound, but make them unsuitable in a streaming model. A note on multiplicative bound is due here. We believe that the general application of LRA is for thin matrices with very small tail singular values (for example, see, [42]). Therefore, we feel that in practical scenario, polynomial multiplicative error would not be that big an issue. The additive error bound computed by [23] $\tau \leq O(k^2\varepsilon^{-1}\sqrt{(\mathsf{rank}(A)\mu + k\log n)\log(1/\delta)\log n})$ compared to $\tau \leq \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}$ (Theorem 4.1(ii)).

Kapralov and Talwar [28]: The only assumption this paper makes is that of singular value separation of the same form as in Hardt and Roth [23]. They also give low-rank approximation in the spectral norm. Additionally, they achieve $(\varepsilon, 0)$-differential privacy, which is only achieved by Chaudhary *et al.* [9]. They use the same definition as in Chaudhary *et al.* [9]. They also sample a singular vector from exponential distribution, but they give a net-based algorithm to perform the sampling in polynomial time.

On the negative side, their mechanism uses k rounds; therefore, it cannot be implemented in a streaming fashion. Since the notion of neighbouring data-sets and privacy guarantee achieved is different from that of ours, we believe our result is incomparable to that of [28]. However, if we just concentrate on the additive error bound, they achieve a bound of $O(dk^3/(\varepsilon\gamma^2))$, where γ is the separation between the singular values, compared to our bound $\tau \leq \frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}$ (Theorem 4.1(ii)).

Hardt and Price [21]: In this recent work, Hardt and Price [21] gave a robust subspace iteration mechanism that allows to publish LRA with noise independent of the rank of the input matrix, thereby, resolving one of the open problems in [23]. They define neighbouring data-sets in the same manner as in [22, 23]. However, they also make an assumption on the singular value separation–a separation between the k-th and $(k+1)$-th singular value of the input matrix. Their mechanism uses k rounds of subspace generation, each of which depends on the spectrum of the matrix to reduce the multiplicative error; therefore, it cannot be implemented in a streaming fashion. We achieve a bound of $\frac{2\sqrt{k(n+d)\ln(2/\delta)}}{\varepsilon}$ (Theorem 4.1(ii)) compared to $\lambda_1\sqrt{kn\mu\log(1/\delta)\log(n/\gamma)\log\log(n/\gamma)}/\varepsilon\gamma^{1.5}\lambda_k$ of [21].

Dwork *et al.* [17]: Dwork *et al.* [17] gave the first single-pass online learning algorithm for private low-rank approximation under the assumption that the rows of the input matrix are normalized. They consider the online-learning model, which is very different from our model, and we do not see any natural way to compare. They use the *follow the perturbed leader* (FTL) algorithm of Kannan and Vempala [26] with the binary tree technique of Dwork *et al.* [13]. This idea was previously used by Jain *et al.* [25] as well. They give a bound that assumes a lower bound of $k\sqrt{n}\log^2(m/\delta)/\varepsilon^2$ on the optimal value, where $\delta < 1/\operatorname{poly}(n)$. More concretely, if OPT is the optimal value, then their error bound is $O(\sqrt{k\mathsf{OPT}}n^{1/4}\log^2(m/\delta))$. We do not make any of the assumptions made by them and, if we just consider the end result, we achieve a bound which is factor $k\sqrt{n}$ better than theirs (see Theorem 4.1(ii)). The case that we are able to bypass their lower bound gives a mathematical indication that the unit norm notion of neighbouring data is strictly weaker than user-level privacy.

□

5 Other Applications of Private Sketch Generation

In this section, we give two applications of PSG_1: mechanisms for MAT-MULT and LIN-REG. Theorem 3.1 guarantees privacy if the singular values are high enough, so our basic approach would be to lift the singular values of the private matrix above the threshold of Theorem 3.1(i). However, we have to be careful. For example, if we use the affine transformation based approach of Blocki *et al.* [4] and Upadhyay [39], then it would lead to an additive error proportional to the Frobenius norm of input matrices. To control the additive error, we follow a different approach. We first transform any conforming matrices $\mathbf{A}$ and $\mathbf{B}$ to $\mathbf{A}'$ and $\mathbf{B}'$, respectively, and then use the identity, $\begin{pmatrix}\mathbb{I} & \mathbf{A}'\end{pmatrix}\begin{pmatrix}\mathbb{I} & \mathbf{B}'\end{pmatrix}^{\mathsf{T}} = \left(\mathbb{I} + \mathbf{A}'\mathbf{B}'^{T}\right)$ to perturb the input matrix with a careful choice of parameters. Intuitively, τ is due to the identity term of the published matrix. We use the same idea for LIN-REG as well.

5.1 Matrix Multiplication

We present the mechanism for MAT-MULT below and the result is stated in Theorem 5.1. The main idea is to lift the spectra of the input matrices above the threshold of Theorem 3.1(i).

INITIALIZATION. On input parameters $\alpha, \beta, \varepsilon, \delta$, set $r = O(\log(1/\beta)/\alpha^2)$. Set $s = \sqrt{16r\ln(\frac{2}{\delta})}/\varepsilon\ln(\frac{16r}{\delta})$. Set the intial sketches of $\mathbf{A}$ and $\mathbf{B}$ to be all zero matrices $\mathbf{Y}_{\mathbf{A}_0}$ and $\mathbf{Y}_{\mathbf{B}_0}$.

DATA-STRUCTURE UPDATE. Set $d = \max\{d_1, d_2\}$. On input a column a of an $n \times d_1$ matrix $\mathbf{A}$ and column b of an $n \times d_2$ matrix $\mathbf{B}$ at time epoch t, set the column vector $\widehat{\mathbf{A}}_{:a} = \begin{pmatrix} s\mathbf{e}_a & \mathbf{0}^{n+d} & \mathbf{A}_{:a} \end{pmatrix}$ and $\widehat{\mathbf{B}}_{:b} = \begin{pmatrix} s\mathbf{e}_b & \mathbf{0}^{n+d} & \mathbf{B}_{:b} \end{pmatrix}$. Invoke PSG_1 with inputs $(\widehat{\mathbf{A}}_{:a}, r, n+d)$ and $(\widehat{\mathbf{B}}_{:b}, r, n+d)$. Update the sketches by replacing the columns a of $\mathbf{Y}_{\mathbf{A}_{t-1}}$ and b of $\mathbf{Y}_{\mathbf{B}_{t-1}}$ by the respective returned sketches to get the sketch $\mathbf{Y}_{\mathbf{A}_t}, \mathbf{Y}_{\mathbf{B}_t}$.

ANSWERING MATRIX PRODUCT. On request to compute the product at time t, compute $\mathbf{Y}_{\mathbf{A}_t}^{\mathsf{T}} \mathbf{Y}_{\mathbf{B}_t}$.

Theorem 5.1. Let $\mathbf{\Omega}$ be the random matrix used by PSG. Then, the data-structure generated by mechanism above uses $O(d\alpha^{-2}\kappa \log(1/\beta))$ bits of space, and on input conforming matrices $\mathbf{A}$ and $\mathbf{B}$, computes (α, β, τ)-MAT-MULT with $\tau \leq s^2\sqrt{n}\alpha$ additive error and (ε, δ)-differential privacy.

Proof. The proof of the utility of mechanism in Section 5.1 follows readily from Lemma 5.2, which is the variance bound computed by [27].

Lemma 5.2. Let $\mathbf{\Omega}$ be a $r \times n$ matrix as constructed in Section 3 with every entries picked from the distribution $\mathcal{N}(0, 1)$, then for a set of m vectors, $\mathbf{v}_1, \cdots, \mathbf{v}_m \in \mathbb{R}^n$, with probability at least $1 - 2\exp(-r\alpha^2/8)$, for any pair $\mathbf{v}_i, \mathbf{v}_j$, we have $|\langle \mathbf{\Omega}\mathbf{v}_i, \mathbf{\Omega}\mathbf{v}_j \rangle - \langle \mathbf{v}_i, \mathbf{v}_j \rangle| \leq \alpha \|\mathbf{v}_i\| \cdot \|\mathbf{v}_j\|$.

The basic intuition is that the multiplication of the scaled identity matrix causes the additive error while the multiplicative error is due to the result of [27]. We follow up with the details. We need to upper bound the quantity $\|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \widehat{\mathbf{A}}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\widehat{\mathbf{B}}/s^2\|_F$. First note that

$$\widehat{\mathbf{A}}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\widehat{\mathbf{B}} = \begin{pmatrix} s\mathbb{I} & 0 & 0 & \mathbf{A}^{\mathsf{T}} \end{pmatrix} \mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega} \begin{pmatrix} s\mathbb{I} & 0 & 0 & \mathbf{B}^{\mathsf{T}} \end{pmatrix}^{\mathsf{T}} = (s^2\mathbb{I}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbb{I} + \mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B}).$$

Therefore,

$$\begin{aligned} \|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \widehat{\mathbf{A}}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\widehat{\mathbf{B}}\|_F &= \|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B} - s^2\mathbb{I}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbb{I}\|_F \\ &\leq \|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B} + s^2\mathbb{I} - s^2\mathbb{I}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbb{I}\|_F \\ &\leq \|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B}\|_F + s^2\|\mathbb{I} - \mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\|_F. \end{aligned} \tag{13}$$

To bound the first term, let random variable X_{ij} denote $(\mathbf{A}^{\mathsf{T}}\mathbf{B})_{ij} - (\mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B})_{ij}$. Then, with probability at least $1 - 2\exp(-k\alpha^2/8)$, we have $|X_{ij}| \leq \alpha\|\mathbf{A}_{:i}\|_2 \cdot \|\mathbf{B}_{:j}\|_2$. Using Lemma 5.2, this results in

$$\|\mathbf{A}^{\mathsf{T}}\mathbf{B} - \mathbf{A}^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{B}\|_F^2 = \sum |X_{ij}|^2 \leq \sum \alpha^2\|\mathbf{A}_{i:}\|_2^2\|\mathbf{B}_{:j}\|_2^2 \leq \alpha^2\|\mathbf{A}\|_F^2\|\mathbf{B}\|_F^2. \tag{14}$$

For the second term, we need to bound the variance on unitaries. This follows from the following set of inequalities.

$$\begin{aligned} \|\mathbf{U}_1^{\mathsf{T}}\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega}\mathbf{U}_2 - \mathbf{U}_1\mathbf{U}_2\|_2 &= \|\mathbf{U}_1(\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega} - \mathbb{I})\mathbf{U}_2\|_2 = \|\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega} - \mathbb{I}\|_2 \\ &= \left(\max \frac{x^{\mathsf{T}}(\mathbf{\Omega}^{\mathsf{T}}\mathbf{\Omega} - \mathbb{I})x}{\langle x, x \rangle} \right) \leq ((1+\alpha) - 1) = \alpha, \end{aligned} \tag{15}$$

where the inequality follows from Theorem 2.6 and noting that r still satisfies the theorem requirement. The result follows by adjusting the value of α after plugging this and equation (14) in equation (13), using the fact that $\|\mathbf{X}\|_2 \leq \|\mathbf{X}\|_F \leq \sqrt{n}\|\mathbf{X}\|_2$ for any $n \times n$ matrix $\mathbf{X}$.

□

5.2 Linear Regression

Our mechanism for LIN-REG is presented below and the result is stated in Theorem 5.3. The main idea is to lift the singular values of the input matrix above the threshold of Theorem 3.1(i).

INITIALIZATION. On input parameters $\alpha, \beta, \varepsilon, \delta$, set $r = O(d\log(1/\beta)/\alpha)$, $s = \sqrt{16r\ln(2/\delta)}\varepsilon^{-1}\ln(16r/\delta)$, and $\mathbf{Y}_{\mathbf{A}_0}$ to be all zero matrix.

DATA-STRUCTURE UPDATE. On input a column c of an $n \times d$ matrix $\mathbf{A}$ at time epoch t, set the column vector $\widehat{\mathbf{A}}_{:c} = \begin{pmatrix} s\mathbf{e}_c & \mathbf{0}^{n+d} & \mathbf{A}_{:c} \end{pmatrix}$. Call PSG_1 with input $(\widehat{\mathbf{A}}_{:c}, r, n+d)$. Update the sketch of $\mathbf{A}$ by replacing the column c of $\mathbf{Y}_{\mathbf{A}_{t-1}}$ by the returned sketch to get the sketch $\mathbf{Y}_{\mathbf{A}_t}$.

ANSWERING QUERIES. On being queried with a vector $\mathbf{b}_i$, set the column vector $\widehat{\mathbf{b}}_i = \begin{pmatrix}\mathbf{0}^d & \mathbf{0}^{n+d} & \mathbf{b}_i\end{pmatrix}$. Call PSG_1 with input $(\widehat{\mathbf{b}}_i, r, n+d)$ to get the sketch $\mathbf{Y}_{\mathbf{b}_i}$. Compute a vector $\mathbf{x}_i$ satisfying $\min_{\mathbf{x}} \|\mathbf{Y}_{\mathbf{A}_t}\mathbf{x}_i - \mathbf{Y}_{\mathbf{b}_i}\|$.

Theorem 5.3. Let $\mathbf{\Omega}$ be $r \times 2(n+d)$ matrix used by PSG, where $r = O(d\log(1/\beta)/\alpha)$. Then the data-structure generated above requires $O(d^2\alpha^{-1}\kappa\log(1/\beta))$ bits and allows to solve (α,β,τ)-LIN-REG problem in an (ε,δ)-differentially private manner with $\tau \leq O(s^2\sqrt{n}\alpha)$.

Proof. The formal proof of Theorem 5.3 is identical to [36], which was refined in [10], modulo the analysis to consider the lift of singular value and using the bound of [27] (Lemma 5.5). We first start with the intuition. For this, let us recall the main result of [10] with matrices $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$. Casted with these matrices, the result of [10] is as follows:

Theorem 5.4. Let $\mathbf{\Omega}$ be as in Section 5.2, $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ be the matrices as constructed in the mechanism stated in Section 5.2. Then with probability at least $1-\beta$, $\|\widehat{\mathbf{A}}\widetilde{\mathbf{X}} - \widehat{\mathbf{B}}\| \leq (1+\alpha)\|\widehat{\mathbf{A}}\widehat{\mathbf{X}} - \widehat{\mathbf{B}}\|$, where

$$\widehat{\mathbf{X}} = \arg\min_{\mathbf{X}} \|\mathbf{\Omega}(\widehat{\mathbf{A}}X - \widehat{\mathbf{B}})\|_F^2 \qquad \widetilde{\mathbf{X}} = \arg\min_{\mathbf{X}} \|(\widehat{\mathbf{A}}X - \widehat{\mathbf{B}})\|_F^2.$$

Now note that $\widehat{\mathbf{A}}^\mathsf{T} = \begin{pmatrix}\mathbb{I} & A\end{pmatrix}^\mathsf{T}$ and $\widehat{\mathbf{B}}^\mathsf{T} = \begin{pmatrix}\mathbb{I} & B\end{pmatrix}^\mathsf{T}$. In other words, both $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ have the same block matrices for the first few rows. Since, $\widetilde{\mathbf{X}}$ minimizes the value of $(\widehat{\mathbf{A}}\mathbf{X} - \widehat{\mathbf{B}})$, it has entries 1 except for the last n positions. Therefore, one way to look at the bound is that the multiplicative error is due to the last n entries of $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ and the additive error is due to the approximation given by the sketch on the rest of the entries.

We follow the approach of [10], which is the refinement of [36]. Their approach to prove the utility bound for linear regression works in two stages and we reiterate it here. We first bound $\|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widetilde{\mathbf{X}} - \widehat{\mathbf{X}})\|_F$, where $\widehat{\mathbf{X}}$ is the right value of the linear regression, $\widetilde{\mathbf{X}}$ is the value of regression for the sketch of $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$, and $\mathbf{U}$ is left singular matrix of $\widehat{\mathbf{A}}$. Let $\widetilde{\mathbf{X}}'$ represents the last n entries of $\widetilde{\mathbf{X}}$ (similarly define for $\widehat{\mathbf{X}}$) and $\widetilde{\mathbf{X}}''$ be the rest of the entries (similarly for $\widehat{\mathbf{X}}$).

The following lemma given by [27] is key to our analysis. This is an equivalent for equation (15) with this value of r. Note that a random Gaussian matrix of the order we considered is a Johnson-Lindenstrauss transform.

Lemma 5.5. Given $r = O(d\log(1/\beta)/\alpha)$. Let $\mathbf{U}$ be any unitary matrix. If $\mathbf{\Omega}$ satisfies the Johnson-Lindenstrauss bound, then with probability at least $1-\beta$, we have $\|\mathbf{U}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\mathbf{U}^\mathsf{T} - \mathbb{I}\|_2 \leq \alpha$.

Using Lemma 5.5, the bound on $\|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widetilde{\mathbf{X}} - \widehat{\mathbf{X}})\|_F$ follows from the triangle inequality. The remainder of the proof follows from the Pythagorus theorem and an observation that A and $\mathbf{U}$ have the same column-space. A structural property from our construction that we repeatedly use for moving from $\widehat{\mathbf{A}}$ to A (respectively, from $\widehat{\mathbf{A}}$ to B) is that except for the block matrix formed by the last n rows, the rest of the entries of the matrices $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{B}}$ are identical. The details follows.

Let B be the matrix formed by the set of queries $\{\mathbf{b}_1, \cdots, \mathbf{b}_m\}$ and $\mathbf{U}$ be the left singular matrix of $\widehat{\mathbf{A}}$. Since the columns of $\mathbf{U}$ is a set of orthonormal vectors, we have $\mathbf{U}\mathbf{U}^\mathsf{T}\mathbf{U} = \mathbf{U}$ and $\|\mathbf{U}^\mathsf{T}\mathbf{U}\mathbf{C}\|_F = \|\mathbf{U}\mathbf{C}\|_F$ for any matrix $\mathbf{C}$. Therefore, it suffices for the utility bound to prove a bound on $\|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|$. For this, we first prove that $\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})$ has a small norm. We have

$$\begin{aligned}\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}}) &= \mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}}) + \mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}(\widehat{\mathbf{B}} - \widehat{\mathbf{A}}\widehat{\mathbf{X}}) \\ &= \mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}(\widehat{\mathbf{B}} - \widehat{\mathbf{A}}\widetilde{\mathbf{X}}).\end{aligned}$$

This is because $\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}(\widehat{\mathbf{A}}\widehat{\mathbf{X}} - \widehat{\mathbf{B}}) = \widehat{\mathbf{A}}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}(\widehat{\mathbf{A}}\widehat{\mathbf{X}} - \widehat{\mathbf{B}}) = 0$. Therefore, from Theorem 5.1 with $\alpha' = \sqrt{\alpha/d}$ (since we chose r in Theorem 5.1 which differs by a factor of α and $1/d$ with respect to that in Theorem 5.3), we have

$$\begin{aligned}\|\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|_F &= \|\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}(\widehat{\mathbf{B}} - \widehat{\mathbf{A}}\widetilde{\mathbf{X}})\|_F \\ &\leq \sqrt{\alpha}\|\mathbf{B} - \mathbf{A}\widetilde{\mathbf{X}}'\|_F + \sqrt{\tau}\end{aligned}$$

From the sub-additivity of the norm and property of conforming matrices (Lemma 2.1), we have

$$\begin{aligned}\|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|_F &\leq \|\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|_F + \|\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}}) - \mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|_F \\ &\leq \sqrt{\alpha}\|\mathbf{B} - \mathbf{A}\widetilde{\mathbf{X}}'\|_F + \sqrt{\tau} + \|\mathbf{U}^\mathsf{T}\mathbf{\Omega}^\mathsf{T}\mathbf{\Omega}U - \mathbb{I}\|_2 \cdot \|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widehat{\mathbf{X}} - \widetilde{\mathbf{X}})\|_F\end{aligned}$$

Using Lemma 5.5 and rearranging the terms, we get $\|\mathbf{U}^\mathsf{T}\widehat{\mathbf{A}}(\widehat{\mathbf{X}}-\widetilde{\mathbf{X}})\|_F \leq 2\sqrt{\alpha}\|\mathbf{B}-\mathbf{A}\widehat{\mathbf{X}}'\|_F + \sqrt{\tau}$. The utility proof is now immediate by observing that the column-space of $\widehat{\mathbf{A}}$ and $\mathbf{U}$ are the same, $\widetilde{\mathbf{X}}'$ is the optimal solution of the regression (which from the normal equations of $\widetilde{\mathbf{X}}$ implies $\mathbf{U}^\mathsf{T}(\widehat{\mathbf{A}}\widetilde{\mathbf{X}}-\widehat{\mathbf{B}}) = \widehat{\mathbf{A}}^\mathsf{T}(\widehat{\mathbf{A}}\widetilde{\mathbf{X}}-\widehat{\mathbf{B}}) = 0$), and the Pythagorus theorem on the norms. More concretely, with probability at least $1-2\beta$

$$\begin{aligned}\|\mathbf{A}\widehat{\mathbf{X}}'-\mathbf{B}\|_F^2 &= \|A\widetilde{\mathbf{X}}'-\mathbf{B}\|_F^2 + \|\mathbf{A}(\widehat{\mathbf{X}}'-\widetilde{\mathbf{X}}')\|_F^2 \\ &\leq (1+4\alpha)\|\mathbf{A}\widetilde{\mathbf{X}}'-\mathbf{B}\|_F + \tau.\end{aligned}$$

Adjusting and renaming the values of α and β, we get the claim of the theorem.

□

LIN-REG IN OTHER MODELS. LIN-REG has been also studied in the local privacy model by Duchi *et al.* [11] and in the online learning model by Jain *et al.* [25] and Thakurata and Smith [38]. These models are different from ours and, therefore, our results are incomparable to theirs. However, if one wishes to make any comparison with our result, we give a brief overview of these results. The readers are welcome to compare it with Theorem 5.3. Jain *et al.* [25] used R queries and T training data set to give a bound of $\tilde{O}\left((R^6\log(1/\delta)\sqrt{n}\log^{1.5}T)/\sqrt{\varepsilon}\alpha^3\right)$. Thakurata and Smith [38] gave a generic algorithm of which Jain *et al.* [25] is a special case. Duchi *et al.* [11] gave a characterization of LIN-REG in the local-privacy model.

5.3 Applications in Learning Theory

Due to lack of space, we do not cover few applications and just make a short note here. For example, PSG can be used in manifold learning. Here, we can consider the streaming version of Baranuik and Wakin [2, Sec 3], where the sampled points are streamed. We first perform a transformation as done in LIN-REG on their sampled points to boost the singular values of the matrix formed by the set of all sampled points and then use Gaussian matrix as the Johnson-Lindenstrauss transform in their algorithm. The more formal treatment of this is done by Upadhyay [40]. Also, our mechanism for LRA can be easily compiled to give differentially private *principal component analysis* using standard algorithms that use LRA on the input matrix as the first step. This achieves almost the same memory bound as achieved by [30].

References

[1] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 20–29. ACM, 1996.

[2] Richard G Baraniuk and Michael B Wakin. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009.

[3] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology–CRYPTO 2008*, pages 451–468. Springer, 2008.

[4] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. In *FOCS*, pages 410–419, 2012.

[5] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *ITCS*, pages 87–96, 2013.

[6] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In Chen Li, editor, *PODS*, pages 128–138. ACM, 2005.

[7] E Bura and R Pfeiffer. On the distribution of the left singular vectors of a random matrix and its applications. *Statistics & Probability Letters*, 78(15):2275–2280, 2008.

[8] Lynn E Cannon. A cellular computer to implement the kalman filter algorithm. Technical report, DTIC Document, 1969.

[9] Kamalika Chaudhuri, Anand D Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *NIPS*, pages 998–1006, 2012.

[10] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *STOC*, pages 205–214, 2009.

[11] John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *Foundations of Computer Science (FOCS)*, pages 429–438. IEEE, 2013.

[12] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In Shai Halevi and Tal Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[13] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *STOC*, pages 715–724, 2010.

[14] Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy N Rothblum, and Sergey Yekhanin. Pan-Private Streaming Algorithms. In *ICS*, pages 66–80, 2010.

[15] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.

[16] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and Differential Privacy. In *FOCS*, pages 51–60. IEEE Computer Society, 2010.

[17] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze Gauss: Optimal Bounds for Privacy-Preserving Principal Component Analysis. In *STOC*, pages 11–20, 2014.

[18] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.

[19] Arjun K Gupta and Daya K Nagar. *Matrix variate distributions*, volume 104. CRC Press, 1999.

[20] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

[21] Moritz Hardt and Eric Price. The noisy power method: A meta algorithm with applications. In *Advances in Neural Information Processing Systems*, pages 2861–2869, 2014.

[22] Moritz Hardt and Aaron Roth. Beating randomized response on incoherent matrices. In *STOC*, pages 1255–1268, 2012.

[23] Moritz Hardt and Aaron Roth. Beyond worst-case analysis in private singular vector computation. In *STOC*, pages 331–340, 2013.

[24] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 604–613, 1998.

[25] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially Private Online Learning. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT*, volume 23 of *JMLR Proceedings*, pages 24.1–24.34. JMLR.org, 2012.

[26] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

[27] Daniel M. Kane and Jelani Nelson. Sparser Johnson-Lindenstrauss Transforms. *J. ACM*, 61(1):4, 2014.

[28] Michael Kapralov and Kunal Talwar. On differentially private low rank approximation. In *SODA*, volume 5, page 1. SIAM, 2013.

[29] Leon Mirsky. Symmetric gauge functions and unitarily invariant norms. *The quarterly journal of mathematics*, 11(1):50–59, 1960.

[30] Ioannis Mitliagkas, Constantine Caramanis, and Prateek Jain. Memory limited, streaming pca. In *Advances in Neural Information Processing Systems*, pages 2886–2894, 2013.

[31] Elchanan Mossel, Ryan O'Donnell, and Rocco P Servedio. Learning juntas. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 206–212. ACM, 2003.

[32] Robb J Muirhead. *Aspects of multivariate statistical theory*, volume 197. John Wiley & Sons, 2009.

[33] C Radhakrishna Rao. *Linear statistical inference and its applications*, volume 22. John Wiley & Sons, 2009.

[34] Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics*, 62(12):1707–1739, 2009.

[35] Youcef Saad and Martin H Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on scientific and statistical computing*, 7(3):856–869, 1986.

[36] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 143–152. IEEE, 2006.

[37] Terence Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.

[38] Abhradeep Guha Thakurta and Adam Smith. (nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems*, pages 2733–2741, 2013.

[39] Jalaj Upadhyay. Random Projections, Graph Sparsification, and Differential Privacy. In *ASIACRYPT (1)*, pages 276–295, 2013.

[40] Jalaj Upadhyay. Randomness Efficient Fast-Johnson-Lindenstrauss Transform with Applications in Differential Privacy and Compressed Sensing . *CoRR*, abs/1410.2470, 2014.

[41] Thomas J Walsh, István Szita, Carlos Diuk, and Michael L Littman. Exploring compact reinforcement-learning representations with linear regression. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 591–598. AUAI Press, 2009.

[42] Christopher Williams and Matthias Seeger. The effect of the input density distribution on kernel-based classifiers. In *Proceedings of the 17th International Conference on Machine Learning*, number EPFL-CONF-161323, pages 1159–1166, 2000.